(How) Do We Teach Emotions?*

Anjali Adukia Matthew Bonci Paula Dastres
Emileigh Harrison Jake Nicoll Teodora Szasz

July 23, 2026

Abstract

Emotional intelligence is a key component of human capital, shaped partially by educational materials. We use machine learning and generative AI to examine emotional content in public-school textbooks and award-winning children's literature. A stark mismatch emerges: while text exposes children to a broad emotional range, images overwhelmingly depict happiness and calm, and "negatively" valenced emotions rarely appear, even when accompanying text describes other emotions. This pattern persists across time, contexts, and identities. Household purchases and library inventories suggest these inputs may be endogenously shaped by consumer demand favoring "positive" cover imagery, implying market forces narrow the emotional landscape children encounter.

Keywords: Culture, Emotions, Curriculum Analysis, Education Policy
JEL Category: I20, I21, I28, J15, Z13

*Affiliation: University of Chicago, University of Wisconsin-Madison, University of Texas at Austin, NBER, IZA. Contact: adukia@uchicago.edu, mbonci@uchicago.edu, dastresgalla@wisc.edu, emileigh.harrison@austin.utexas.edu, jwnicoll@uchicago.edu, dorha.szasz@gmail.com. For helpful feedback, we thank Alex Eble, James Evans, Yuan Chang Leong, Lesley Turner, Isabeau Dasho, and seminar participants at Notre Dame, Zurich AI + Economics, Hong Kong AI Research in the Social Sciences and Humanities, Association for Education Finance and Policy, the Society for Research on Educational Effectiveness, Stanford, and The University of Chicago. For technical and research assistance, we thank Ping-Jung Liu. For financial support, we thank the Becker Friedman Institute at UChicago. The research reported here was also supported by the Institute of Education Sciences, U.S. Department of Education, through Grant R305A200478 to the University of Chicago. The opinions expressed are those of the authors and do not represent views of the Institute or the U.S. Department of Education. Researcher(s)' own analyses calculated (or derived) based in part on data from Market Track, LLC dba Numerator and marketing databases provided through the Numerator Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Numerator data are those of the researcher(s) and do not reflect the views of Numerator. Numerator is not responsible for and had no role in analyzing or preparing the results reported herein.

The centrality and universality of emotion to human experience has long been recognized across cultures, traditions, and disciplines (Izard, 1977; Lutz and White, 1986; Elster, 1998; Marcus, 2000; Gross, 2015; Denham, 2019). Beyond their psychological importance, emotional skills constitute a key component of human capital. Individuals with greater emotional intelligence exhibit stronger stress management, prosocial behavior, academic performance, and workplace functioning (Mayer and Geher, 1996; Parker et al., 2004; Austin, Saklofske and Mastoras, 2010; Brackett et al., 2012; Angelucci, Fábregas and Vazquez, 2026). This may be especially important in modern labor markets; as work has become more team-based and service-oriented, employers increasingly seek workers who can communicate effectively, respond to conflict, and manage interpersonal relations (Hochschild, 1983; Deming, 2017; Bonhomme, 2021; Hansen et al., 2021; Hoffman and Tadelis, 2021).

Despite growing evidence on the returns to social and emotional skills, less is known about the childhood inputs that develop these skills. We therefore explore an upstream question: where do emotional skills come from and, specifically, what are children taught to recognize or display? Standard models of human-capital formation emphasize the importance of early investments such as those made through schools. Books and other children's media serve as one such site of learning. The content of those books featured in school curricula and on library displays legitimize and institutionalize particular knowledge, values, and norms. As the primary instructional resources in classrooms, they not only transmit academic content, but they also signal to students – through both explicit content and implicit messaging – which competencies, perspectives, and emotional dispositions are deemed socially and economically valued.

In this study, we examine emotions children encounter in educational materials across formal and informal settings. We first analyze the emotions presented across contexts, identities, and over time using two sets of influential children's media: (i) three decades of public-school textbooks from California and Texas and (ii) a century of award-winning children's books commonly found in homes, libraries, and schools (Smith, 2013; Cockcroft, 2018; Adukia et al., 2023). We then incorporate data on both individual and library purchasing behavior to understand the potential role of market forces in shaping the presence of particular emotional portrayals.

To extract data on emotions from books, we apply artificial intelligence methods from natural language processing (applied to text) and computer vision (applied to images). Specifically, we employ large multimodal models (LMMs) to classify eight primary emotions across both text and images: anger, calm, confusion, disgust, fear, happiness, sadness, and surprise.[1] To assess classification robustness, we compare outputs from our primary models with alternative approaches, including a lexicon-based emotion classification method applied to text, a RoBERTa transformer

[1]This set of emotions was derived from the six basic emotions identified by Ekman, Sorenson and Friesen (1969) with two added emotions: calm and confusion, particularly relevant to children's media contexts. Given that we are predicting these emotions in the text and images of children's media, we also allow for the model to predict "unsure" or "no emotion" to allow for cases that are not obvious or cases where a character is detected but their emotion is not visible.

model fine-tuned on an emotion-specific text dataset, and a pre-trained convolutional neural network (CNN) used for feature-based emotion classification in images. We examine heterogeneity by context (e.g. corpora, state, grade, subject) and character identities (e.g. putative race, skin color, gender, age, human).

Our analysis reveals a systematic divergence between textual and visual emotional content: written text exposes children to a broad emotional range including happiness, sadness, anger, and fear. By contrast, images almost exclusively depict happiness and calm while rarely representing emotions traditionally categorized as "negative." These patterns persist across time periods, contexts, and demographic groups both in text and in images. Moreover, we find a striking misalignment between the emotions portrayed in the text and images on the same page. Regardless of which emotions characters display in the text, images disproportionately depict characters as calm or happy. This demonstrates that the discrepancy between emotional representation in text and images does not simply reflect choices by illustrators to highlight certain positive components of their narratives; rather, the disconnect reflects a genuine incongruence between the emotions that are written about and the emotions that are shown. Almost uniformly, these books suggest that regardless of the emotions characters experience in the text, the appropriate outward response is to display a calm or happy appearance.

We next explore the role of consumer demand in the visual overrepresentation of calm and happiness in children's books. By merging information about the emotions in children's books with purchaser decisions – both of households and public libraries – we find that, controlling for book attributes, purchasers are less likely to choose children's books whose covers depict emotions traditionally considered "negative." This suggests that the lack of a range of emotions in images may be endogenously shaped by purchaser demand and reflects the societal value placed on displays of "positive" emotions. This may be because these emotions capture attention and make books look more appealing (Bloch, 1995; Underwood and Klein, 2002; Pieters and Wedel, 2004). Additionally, if purchasers have limited bandwidth to accurately assess the quality of emotional representation in books and the messages they convey, then authors and illustrators could be incentivized to oversupply happy and calm images in equilibrium even relative to the underlying preference for these emotions (Ellison, 2006; Cusumano, Fabbri and Pieroth, 2024).

Our findings contribute to multiple strands of literature. First, we contribute to our understanding of the production function of socioemotional skill by exploring a source of how children are taught to recognize and display emotions. This informs the literature on non-cognitive skill development (Groves, 2005; Heckman, Stixrud and Urzua, 2006; Duckworth et al., 2007; Kautz et al., 2014; Lundberg, 2017), highlighting how curricular materials can shape early emotional intelligence. In doing so, we transpose Elster's (1998) argument that emotions themselves matter for choice. Instead of asking how emotions affect adult economic decisions, we ask how emotional display norms may be produced through childhood exposure. This extends the idea beyond emotions being not only inputs into economic behavior; they are themselves shaped by institutions, markets, and cultural

goods.

Second, we add to the literature on cultural economics, which has recently rapidly expanded through the use of novel data sources to study naming conventions (Bazzi, Fiszbein and Gebresilasse, 2020), folklore (Michalopoulos and Xue, 2021), movies (Michalopoulos and Rauh, 2024), style (Yanagizawa-Drott and Voth, 2023), and now emotions. We extend this agenda from *who* is represented to how people are emotionally represented. The core object is an emotional norm: the text describes many emotions to children, but visually children primarily see calm and happiness, which suggests there are "emotional display rules" embedded in culture.

Third, we expand prior research on emotional representation in children's books (Nikolajeva, 2014; Ku, 2019; Kardum, Dadic and Horvat, 2021), which has largely relied on manual content analysis and focused narrowly on text. This analysis helps to understand the socioemotional curriculum embedded in children's content in both formal and informal settings.

Fourth, we contribute to the economics of media and product market literature, where content is not treated as exogenous but rather firms choose content in response to consumer demand (Gentzkow and Shapiro, 2010). We show these patterns may be equilibrium outcomes of markets for children's media. Moreover, we illustrate how books can effectively advertise themselves to purchasers with limited attention to a book's content (Dröge and Darmon, 1987; Eliaz and Spiegler, 2011; Bordalo, Gennaioli and Shleifer, 2013), in this case, through the emphasis of certain emotions on covers, which are salient, easy-to-process signals. Indeed, private purchasing decisions may generate an equilibrium with unrepresentative visual emotional diversity for children's books.

Finally, we contribute to computational social science by developing a multimodal approach to measuring emotional representation in children's media, and in particular showing misalignment in messages across modalities. Building on work that uses natural language processing and computer vision to study social representation, stereotypes, and cultural content (Ash et al., 2021; Adukia et al., 2023; Ludwig and Mullainathan, 2024; Adukia and Harrison, 2025), we examine the emotional content in the text and images of children's media: public-school textbooks from the last three decades and influential books from a century of children's literature.

The paper proceeds as follows: Section I discusses background related to the study and teaching of emotions. Section II details our data sources. Section III explains the processes used to transform the images and text into data, along with the tools employed for analysis. The main results are presented in Section IV. In Section IV.D, we further probe the apparent contrast between emotions in text and images by considering both modalities on the same page simultaneously. In Section V.A, we demonstrate robustness to alternative approaches to measurement, while Section V.B discusses our model validation strategy, which we hope offers guidance for others using LMMs for measurement in economics and in other disciplines. In Section VI, we explore how consumer demand may contribute to the patterns we uncover. Finally, we offer a discussion in Section VII.

I Background

Emotions constitute a fundamental dimension of human experience, making it a natural subject of sustained inquiry across a wide range of disciplines. In psychology, foundational work has examined the nature, classification, and development of emotions (Izard, 1977; Plutchik and Kellerman, 1980; Averill, 1983), while more recent research has emphasized emotion regulation, emotional competence, preferences over affective states, the influence of emotions on economic and social behavior, and the role of early socialization in shaping how children learn to recognize, manage, and express emotions (Eisenberg, Cumberland and Spinrad, 1998; Gross and John, 2003; Tsai, Knutson and Fung, 2006; Tsai, 2007; Côté, Gyurak and Levenson, 2010; Genevsky and Knutson, 2015; Gross, 2015; Denham, 2019; Park et al., 2020; Perry et al., 2020). Sociology has highlighted how social roles, status relations, institutions, and labor markets shape emotional expression, including norms regarding which emotions are appropriate to display in different contexts (Kemper, 1978; Hochschild, 1983; Thoits, 1989). Anthropological research has emphasized that emotional experiences and displays reflect culturally shared meanings and expectations, including norms governing appropriate emotional responses across social settings (Lutz and White, 1986). Political science has examined the role of emotions in political judgment, participation, and collective behavior (Marcus, 2000). Philosophy has long debated the nature, moral significance, and rationality of emotions (Solomon, 1976; Rorty et al., 1980), while economics has examined how emotions influence preferences, decision-making and social interactions (Elster, 1998; Loewenstein, 2000).

A large literature shows that non-cognitive skills shape educational attainment, employment, and earnings (Heckman and Rubinstein, 2001; Groves, 2005; Heckman, Stixrud and Urzua, 2006). Emotional intelligence – the capacity to recognize, understand, and regulate emotions in oneself and others – is one such skill. For example, manipulating emotional displays to exhibit happiness and calmness is one feature of what customer service workers regularly provide (Hochschild, 1983). Additionally, regulating one's own emotions and recognizing emotions in others can reduce frictions in teams and contribute to managerial effectiveness, which in turn can improve earnings (Hansen et al., 2021; Hoffman and Tadelis, 2021). Consequently, the return to social-emotional skills has risen even as the return to some cognitive skills has stagnated (Castex and Kogan Dechter, 2014; Edin et al., 2022).

Schools are a central institution through which children can acquire and practice emotional skills. In particular, these institutions foster emotional intelligence through direct instruction (Brackett et al., 2012), classroom interactions, and teacher expectations. Education materials and children's books are another important site of emotional display learning (Green and Sun, 2025). According to the constructionist hypothesis in developmental psychology, emotions are abstract concepts learned by children through an accumulation of encounters with emotions (Hoemann, Xu and Barrett, 2019). From this perspective, the images in children's books can provide a model for the accompanying emotion language that can help children calibrate both their emotion recognition skills and their understanding of how to express emotions during the age range that these

skills are developing. Indeed, children learn to recognize emotions between the ages of two and eight, although the skill is still malleable after the age of eight (Widen and Russell, 2008; Ruffman et al., 2023). This calibration could improve emotion recognition skills, which have been shown to translate to improved social relationships (Brackett et al., 2012). Consistent with this hypothesis, research suggests that exposure to emotionally diverse content and educational television programs improves children's emotional vocabulary, perspective-taking skills, empathy, emotion recognition, and coping skills (Kumschick et al., 2014; Rasmussen et al., 2016; Foulds, 2023).

Moreover, visual depictions of emotions may have greater influence than textual presentations of emotions on children's emotion identification and response capabilities (Arizpe and Styles, 2015; Nikolajeva, 2017). This ability to identify emotions may have important social consequences: children who struggle to identify emotions such as sadness tend to exhibit lower prosocial behavior (Doherty, 2008; Dickerson and Quas, 2021). Formal and informal educational materials may play a crucial role in this developmental process.

Because emotional norms are relevant for later educational and labor market success, the way educational materials represent emotions is potentially important. Ex ante, however, it is not obvious which emotions these materials will emphasize. On the one hand, educational materials may aim to build emotional intelligence by exposing children to the complexity of real-life emotional experience. Consistent with the constructionist hypothesis, this could involve stories and images that portray a broad range of emotions and model how characters recognize, express, and regulate them in healthy ways (Hoemann, Xu and Barrett, 2019). On the other hand, educational materials may also reflect prevailing social and labor market norms about which emotions are appropriate to display. Many social situations and occupations discourage the expression of so-called "negative" emotions and instead reward calmness, compliance, or a façade of positivity (Hochschild, 1983; Deming, 2017). To the extent that schools socialize children into future social and labor market roles, these norms may be reflected in materials that downplay or omit emotions such as sadness, anger, or fear (Bowles and Gintis, 1976; Entorf and Dohmen, 2025).

This ambiguity creates an empirical question: do children's educational materials expose children to a broad emotional repertoire, or do they reproduce a narrower set of socially and economically rewarded emotional displays? Moreover, because children encounter these materials through both text and images, it is important to examine not only which emotions are described, but also which emotions are visually modeled. If images do not align with the emotions described in the accompanying text, materials may provide weaker or inconsistent cues for developing emotion recognition skills (Brady, Ogren and Johnson, 2024). Despite the importance of children's materials in the formation of emotional skills, there is limited systematic evidence on how emotions are represented across both modalities.

We address this gap by examining the emotional content of both textbooks and children's literature, comparing patterns across formal and informal educational materials. We also incorporate data on purchaser and public library purchasing behavior to assess whether demand for

particular emotional portrayals influences book acquisition decisions and, in turn, the emotional representations children are likely to encounter.

II Data Sources

We analyze two categories of children's media representing both formal and informal educational contexts: public-school textbooks and award-winning children's books. These materials constitute two prominent sources of content to which children are regularly exposed during critical developmental periods.

For the public-school curricular materials, we analyze state-adopted textbooks from Texas and California in the subjects of Science, Social Studies, and Reading/Language Arts for 3rd and 5th grades. These states are meaningful for several reasons. First, California and Texas are the two largest public-school systems in the United States, enrolling approximately 5.92 million and 5.53 million students, respectively, in fall 2023. Second, California and Texas both use state-level review processes in which appointed reviewers evaluate publisher materials for standards alignment, whereas most states leave textbook selection primarily to local districts (Doan and Kaufman, 2024). Third, these states drive textbook purchasing markets in other states. Finally, these states occupy opposing positions on the political spectrum with Texas advisory panels being more likely to reflect Republican perspectives and California panels being more likely to reflect Democratic perspectives. Differences across states could reflect the influence of ideology on the representation of emotions and identities in educational materials (Biasi et al., 2025).

For children's literature, we focus on books that are likely to be particularly influential. Specifically, we include titles that have won awards featured by the Association for Library Service to Children. These awards honor books for their "exceptional literary" or "artistic" value, or for centering the experiences of people with underrepresented identities. We analyze two sets of children's books: those that have won "Mainstream" awards, which emphasize literary quality and influence, and those recognized with "Diversity" awards, which specifically celebrate books featuring underrepresented identities. The corpus encompasses both fiction and nonfiction works, representing realistic and imaginary settings, thereby capturing the full spectrum of narratives available to young readers in curated, high-quality literature.

III Methods

Previous research has established robust computational methods for emotion detection in textual data (Bollen, Mao and Pepe, 2011; Thelwall, Buckley and Paltoglou, 2011; Hasan, Rundensteiner and Agu, 2014; Demszky et al., 2020). However, emotions are communicated not only through language, but also through non-verbal cues such as gestures and body language. Facial expressions in particular have been extensively documented as reliable indicators of emotional states (Sullivan and Masters, 1988; Lucas, Diener and Larsen, 2003; Öhman, 2007, 2009), so only measuring emotional representation in text may miss a central channel through which children encounter and learn about emotions. We therefore develop a multimodal framework that systematically quantifies

emotional content in both the textual and visual components of children's media.

We focus our analysis on eight emotion categories: anger, calm, confusion, disgust, fear, happiness, sadness, and surprise. These classifications incorporate the six basic emotions originally delineated by Ekman, Sorenson and Friesen (1969) while extending the framework to include calm and confusion. To accommodate ambiguous cases, we include two additional classifications: "unsure" for instances where multiple emotions appear equally plausible or emotional expression is ambiguous and "no emotion" for affectively neutral content, where a character's emotion is not obvious or visible. While the concept of universal basic emotions remains contested in the literature (Russell, 2003), empirical evidence supports the existence of discrete emotional categories (Cowen and Keltner, 2020). However, the precise set of discrete emotions is not agreed upon (Ekman, 2016), and our primary set aligns with taxonomies accounting for displays of emotions across different contexts (Cowen and Keltner, 2017), including in faces (Cowen and Keltner, 2020).[2]

III.A Text Analysis

We developed three complementary approaches to quantify emotional content in the text of children's printed content. Our primary measure uses a large language model (LLM) for character-level emotion detection. We supplement this with a lexicon-based frequency analysis and a fine-tuned transformer model. This multi-method strategy allows us to test the sensitivity of our findings to different measurement strategies. Figure 1a provides an overview of the data extraction process applied to the text on book pages.

Large Language Model Classification. We use a large language model (LLM) (specifically GPT-4o-mini) to construct page-level measures of character emotions in text, motivated by recent evidence that LLMs can support scalable social-science annotation tasks, including emotion classification (Ziems et al., 2024). For each page, the model identifies characters, assigns each character one of eight emotion labels – anger, calm, confusion, disgust, fear, happiness, sadness, or surprise – and records "unsure" or "no emotion" when appropriate. We also extract character attributes, including age, gender, race, and whether the character is human, allowing us to analyze heterogeneity in emotional representation across demographic groups. More details are found in Appendix C.A.

Conceptually, GPT-4o-mini processes the raw text as input and generates character-emotion pairings by: (1) recognizing linguistic markers of emotional expression (e.g., descriptive verbs, dialogue, narrative commentary); (2) matching these markers to semantic patterns learned during model pre-training; and (3) mapping identified emotional content to our taxonomy. The model essentially performs linguistic pattern recognition at the character level, converting narrative de-

[2]Indeed, what constitutes an emotion and whether calmness is an emotion or a state of no emotion has been of debate for quite some time, particularly starting when James (1884) first suggested that emotions to be considered are those with a "distinct bodily expression" followed by contention by Gurney (1884) who noted that even when emotion's "*mass*, and its character *quâ painful*, remain unaffected under these latter conditions, it is just as truly emotion as before, and emotion to which the physical signs...do not now even contribute." This debate has extended throughout time and disciplines (c.f. discussion by philosophers such as Zemach 2001 and neuroscientists such as Adolphs, Mlodinow and Barrett 2019.)

scriptions of emotional states into discrete categorical labels.

This differs from traditional lexicon-based approaches (which count emotion-related keywords and which we employ in subsequent analysis) by leveraging contextual understanding; the model interprets emotional expressions in relation to surrounding narrative, character relationships, and plot development.

Lexicon-Based Analysis. We also implement a lexicon-based approach which involves counting the frequency of predetermined emotion-related words, offering a straightforward and interpretable baseline. To do so, we construct comprehensive keyword dictionaries.[3] For each sentence, we enumerate emotion-specific keyword occurrences and aggregate counts to our analytical units of interest (e.g., complete corpus, subject matter categories, demographic subgroups). This frequency-based approach provides a transparent, replicable baseline measurement independent of neural network architectures.

Transformer Model Prediction. As a second text-based robustness check, we apply a RoBERTa-base transformer model fine-tuned on the GoEmotions dataset created by Demszky et al. (2020).[4] The model returns a probability score for each of the 28 GoEmotions labels, which include 27 emotion labels and a neutral category. In our implementation, we assign each sentence to the label with the highest predicted probability and drop sentences classified as "neutral," since our goal is to characterize the distribution of emotional content.[5] We then map the remaining fine-grained GoEmotions labels to the corresponding aggregate emotion categories used in our analysis.[6] This supervised transformer-based approach provides an additional benchmark for our primary LLM-based text classification method.

III.B Image Analysis

Social scientists have increasingly utilized images as primary sources of information (Ash et al., 2021; Adukia et al., 2023; Yanagizawa-Drott and Voth, 2023; Adukia and Harrison, 2025). Building on this emerging methodological literature, we developed two complementary computational pipelines to systematically extract structured data on emotional expressions and character attributes from images in children's books. We use an LMM and, separately, a two-step face detection model and feature classification model. Figure 1b provides an overview of the data extraction process applied to the images on book pages.

[3] We construct these dictionaries using three sources: WordNet semantic expansion, ChatGPT-assisted generation, and manual curation; more details are found in Appendix C.B. Examples of keywords related to each emotion can be found in Appendix Table B.1. The full lexicon can be viewed at https://github.com/miielab/emotions_paper.

[4] The model can be found at https://huggingface.co/SamLowe/roberta-base-go_emotions.

[5] Dropping neutral sentences parallels our lexicon-based approach, in which we only examine sentences containing emotion-related words.

[6] Because GoEmotions does not include a "calm" label, this approach maps the model outputs to the remaining aggregate emotion categories. We view this limitation as useful for assessing whether our conclusions are sensitive to the set of emotions available to the text classifier. Our mapping is adapted from the correspondence categories provided at https://github.com/google-research/google-research/blob/master/goemotions/data/ekman_mapping.json.

Large Multimodal Model Classification. We use a large multimodal model (LMM) (specifically Gemini 2.5 Pro Preview) to construct character-level measures of emotions and attributes from book images. We process scanned pages in two stages: first, we classify each page as having only text or containing an image; only pages with images proceed to the second stage. For each illustrated page, the model provides information for the illustration as a whole such as a description of the setting as well as information on each depicted character including bounding-box coordinates, race, gender, age, skin color, human/non-human, and emotion. Mirroring our text analysis, our primary emotion taxonomy includes anger, calm, confusion, disgust, fear, happiness, sadness, and surprise, with "unsure" and "no emotion" allowed for ambiguous or neutral cases. More information is provided in Appendix C.B.

Conceptually, Gemini 2.5 Pro Preview analyzes images by: (1) identifying human and non-human faces and figures; (2) extracting visual features associated with emotional expression (e.g. facial muscle configurations, body posture, spatial positioning); (3) mapping these visual patterns to learned semantic categories; and (4) simultaneously inferring demographic characteristics based on visual cues (e.g. phenotypical features).

The model leverages pre-training on a large corpora of annotated images, allowing it to recognize associations between visual features and emotional states. For example, an upturned mouth with engaged eye muscles (as referred to by Ekman (1989) as Duchenne markers of smiling based on Duchenne de Boulogne (1990)) would be classified as "happiness," while a downturned mouth or furrowed brow would map to negative valences. It is important to note that this process is fundamentally visual pattern recognition and should not be construed as interpretation. The model identifies statistical regularities between pixel configurations and emotion labels learned during pre-training, rather than "understanding" emotions as subjective experiences.

Convolutional Neural Network-Based Face Detection and Feature Classification. As a second approach to measuring emotions and character attributes in images, we use a computer-vision pipeline that builds on the face-detection and feature-classification models developed in Szasz et al. (2022). We first apply FDAI, a face-detection model trained on illustrated faces from children's books, to identify faces in book images. We then use an extended version of our feature-classification model, originally developed for Szasz et al. (2022), to classify race, gender, age, and emotion for each detected face. We use these predicted labels as an alternative to our LMM-based classifications to ensure that our results are robust to different measurement choices.

IV Results

In this section, we present our findings on the emotional content expressed through the text and images of influential children's materials.

IV.A Emotions Overall and Over Time

We first examine the overall patterns of emotions in text and in images, respectively. Figure 2 shows the distributions of detected emotions in text and images across all books in our sample. In the

text, the most common emotions are happiness and sadness, which account for approximately 29% and 20% of characters' emotions, respectively (Figure 2a). With respect to the distribution of the other emotions (anger, calm, confusion, disgust, fear, and surprise), we observe a meaningful number of depictions of each in text (ranging between 4% and 12%), reflecting a balance between "positively" and "negatively" valenced emotions. In sharp contrast, there is a high degree of homogeneity in the emotions displayed in images. In images, calm and happiness together account for approximately 90% of the depicted emotions by characters (Figure 2b).[7] Notably, negatively valenced emotions, which are meaningfully present in the text, are largely absent from images. This pattern illuminates an incongruence between the emotional narratives in the text and the limited emotional range portrayed in the images. Figure 3 shows that these findings are stable over time.

IV.B Emotions in Different Contexts

To better understand how emotions are represented across different types of books, we examine heterogeneity across a variety of different subcollections of our data. Overall, we find similar patterns across each context: images remain concentrated in calm and happiness, whereas a wider variety of emotions is displayed in text. We discuss them further here.

We first examine the patterns across different book types. In Figure 4 we see that the distribution of emotions in the text of children's literature and textbooks are broadly similar, with textbooks showing slightly more representation of happiness compared to children's literature. Children's literature displays a somewhat more diverse range of emotions compared to textbooks, though both collections cover a variety of emotions in text.

Within the subcollections of these two sources of children's text, the patterns remain qualitatively similar. In particular, the Mainstream and Diversity subcollections of children's literature show comparable proportions of each emotion. The primary difference is that the Diversity subcollection exhibits a slightly greater proportion of sadness (approximately 5 percentage points higher). With respect to textbooks, we observe similar patterns across those from California and Texas.

While the overall patterns hold across grade levels, we observe noteworthy variation. For example, happiness appears ten percentage points more frequently in third-grade textbooks than fifth-grade textbooks, while emotions such as anger, fear, and sadness appear slightly more frequently in fifth-grade textbooks. This may reflect curricular design priorities with earlier grades emphasizing reassurance and "positive" affect and upper elementary instruction introducing more complex narratives that require a broader emotional range.

Subject matter appears to slightly shape emotional representation. Science textbooks contain fewer depictions of anger and sadness compared to Reading and Social Studies textbooks. This

[7] Note that these proportions reflect the proportion of all characters that have an emotion. We drop characters that are classified as having "No Emotion" or an "Unsure" Emotion. The proportion of pages that fit these categories are approximately 30% for characters in the text and 27% for characters in the images. Distributions including these classifications can be found in Figure A.1. Dropping these characters does not change our results in any substantive way and improves the interpretability of our analyses.

likely reflects disciplinary differences in pedagogical approach for each subject: Science instruction emphasizes concept mastery through factual exposition and may not require as multifarious an emotional narrative, whereas Reading and Social Studies inherently rely on narrative and human experience, which naturally incorporate emotional complexity. Even acknowledging small differences across subjects, we still observe that the representation of emotions in text remains broadly consistent across all subcollections.

Image heterogeneity analysis corroborates the overall image analysis: calm and happiness dominate emotional depictions across all subcollections (Figure 5). There are a few notable patterns that emerge. First, textbooks show higher concentrations of calm and happiness compared to children's literature, which is consistent with the textual analysis and may be due to the nature of children's literature expressing a richer display of distinct emotions compared to school textbooks. Second, this effect varies by subject. Science and Social Studies textbooks show the highest concentration of "positive" emotions, while Reading textbooks display more moderate patterns that more closely resemble children's literature. This suggests that textbooks organized around narrative content (Reading) adopt market conventions from children's publishing. Finally, while grade-level differences in emotional composition emerge (grade 3 favors happiness; grade 5 favors calm), these differences are compositional rather than substantive in valence: the total proportion of positive emotions remains stable across grades.

IV.C Emotions Displayed by Different Identities

Having established the aggregate text-image mismatch, we now examine whether emotional representation varies systematically by character identity. This analysis serves two purposes: it tests the sensitivity of our main findings across demographic subgroups, and it reveals whether educational materials encode different emotional norms for characters of different identities (e.g. genders, ages, races, and human/non-human status). As described in Section III, all identity attributes are model-generated predictions derived from textual description for the text analysis and from visual features for the image analysis.

Figure 6 presents the heterogeneity analysis in text. First, gendered differences emerge: female characters display sadness somewhat more frequently than male characters, while male characters more often display anger. This asymmetry reflects longstanding cultural associations linking women with vulnerability and men with assertiveness or aggression. Second, age-based differences appear: adult characters more frequently express anger and calm, whereas child characters more often express fear and confusion. This pattern may reflect children in narratives being more likely to be positioned as encountering unfamiliar or threatening situations, while adults are portrayed as composed authority figures or, alternatively, as sources of conflict. Importantly, these differences are modest in magnitude. Across all identity groups, textual content preserves substantial emotional diversity: every subgroup displays a meaningful range of emotions rather than a single dominant category.

Figure 7 presents the analogous heterogeneity analysis for images. The near-universal dominance of calm and happiness persists across every identity group, i.e. no subgroup displays negative emotions with meaningful frequency. This pattern applies uniformly to characters of all genders, ages, races, skin color, and human/non-human status. Within this constrained emotional space, however, subtle identity-based patterns emerge in how the two dominant emotions are allocated. Male faces are more likely than female faces to be depicted as calm, while female faces are correspondingly more likely to be depicted as happy: a difference of similar magnitude. This trade-off is consistent with a gendered emotional stereotype in which composure or emotional restraint is coded as masculine, while overt positive expressiveness (smiling) is coded as feminine. A parallel structure holds across other identity groups: subgroups depicted more frequently as calm are depicted less frequently as happy, and vice versa. This near-perfect substitution between calm and happiness underscores the central visual constraint: characters are almost always shown expressing one of these two positive emotions, and identity primarily shapes *which* positive emotion is displayed rather than whether the character is permitted to display a negative one.

IV.D Concurrent Text and Image Analysis

The preceding sections document a mismatch between the aggregate distributions of emotions in text and images. This section sharpens that finding by comparing emotions within the same page. If images simply emphasized the positive moments of otherwise emotionally diverse narratives, we would expect visual and textual emotions to align when both are present. Instead, we find systematic incongruence: images depict calm and happiness even when the accompanying text describes very different emotions.

To further investigate our text and image data, we directly compare the emotions expressed through each modality on the same page. Figure 8 presents the co-occurrence of emotions for text and images on the same page, normalized by the text columns. In other words, each cell reports the share of a specific emotion label in images, given a particular emotion label in the text of the same page. The results reveal a striking asymmetry. Calm is the only emotion that consistently aligns across modalities: when calm appears in the text, 70% of image emotions on that page are also calm. Happiness follows calm, but it aligns more weakly: when the text describes happiness, only 27% of images depict happiness. For every other emotion, alignment collapses. When the text conveys anger, fear, sadness, or surprise, the corresponding images depict that same emotion only 13%, 9%, 10%, and 8% of the time, respectively.

Crucially, the misalignment is directional. When the text expresses an emotion other than calm, images on that page still depict calm between 58% and 67% of the time. In other words, negative textual emotions are systematically accompanied by positive visual expressions. Children reading these pages encounter a consistent implicit message: regardless of what characters say, do, or feel in the narrative, the appropriate outward display is calm or happy.

One concern is that page-level co-occurrence may conflate different characters. For example,

the text might describe one character's anger while the image portrays a different, calm character. To probe this possibility, we recompute the co-occurrence matrix on a subsample of pages where we can identify a unique character of the same predicted race and gender in both the text and image of a particular page. Figure A.3 illustrates that the distributions of emotions in text and images in this matched sample are similar to the overall distributions displayed in Figure 2, indicating that the matched sample is similar to the full sample of text and images.

Figure A.4 then shows the misalignment persists even for matched characters: the diagonal carries similar weight to Figure 8, meaning that a single character's textual and visual emotions frequently diverge. Regardless of the emotion a character displays in the text, that same character is most likely to be depicted as calm in the accompanying image.

We further quantify the overall text-image discordance by computing a page-level measure of the distance between the set of emotions in text and the set of emotions in images. In particular, for each page, we build an eight-dimensional indicator vector separately for text and images, where each entry in these vectors corresponds to an emotion and equals one if and only if the corresponding emotion is present on the page. After normalizing, we compute the total variation distance (TVD) between the page-level text and image distributions, $\delta(p, q) = \frac{1}{2}\|p - q\|_1$. Figure A.5 displays the distribution of these total variation distances. The result is bimodal: in approximately 10% of cases, the TVD is zero, indicating perfect alignment between textual and visual emotions. For nearly half of the pages, however, the TVD equals one, indicating that there is no overlap between the emotions conveyed in text and those conveyed in images on the same page.

Together, these results establish that the text-image mismatch is not an artifact of aggregation or of illustrators selectively highlighting positive narrative moments. The discrepancy holds within pages, within matched characters, and across the full distribution of pages. Illustrations convey happiness and calm even when these emotions are entirely absent from the surrounding text. These patterns speak directly to the mechanisms discussed in Section I. Given the documented role of visual cues in children's emotion recognition, this systematic incongruence may narrow the range of emotional displays modeled to children, may confuse the learning signal that they receive where the emotional expression contradicts the narrative context, and may reinforce an implicit norm that negative emotions, however they are experienced, should be outwardly suppressed.

V Robustness and Validation

We check robustness and validate our approach using alternative methodologies: (1) for text, a hand-curated lexicon-based approach counting emotion-related keywords; (2) for text, a fine-tuned transformer model (RoBERTa) trained on the GoEmotions dataset; (3) for images, a two-stage CNN-based pipeline involving character detection and feature classification; and (4) human annotators on a stratified sample. The consistency of results across these methods substantiates our primary generative-AI-based findings and addresses concerns about model-specific artifacts.

V.A Alternative Approaches to Measurement

Our primary results rely on large multimodal models: GPT-4o-mini for text and Gemini 2.5 Pro Preview for images. A natural concern is that our findings reflect idiosyncrasies of these particular models rather than genuine features of the underlying materials. To address this, we replicate our core analysis using three alternative classifiers: a lexicon-based word-count approach and a fine-tuned transformer (RoBERTa/GoEmotions) for text, and a convolutional neural network (CNN)-based image classifier, described in Sections III.A and III.B. These methods span a wide range of complexity and transparency, from a fully interpretable keyword dictionary to less transparent neural architectures, providing a ranging set of robustness exercises.

we note that these alternatives differ from our primary approach in their unit of analysis, so perfect agreement is neither expected nor desired. The CNN classifier assigns emotions to individual detected faces, whereas the LMM analyzes full illustrations and the characters within them. Similarly, the lexicon and transformer methods operate at the word and sentence level, respectively, while our primary text model produces character-level labels. What matters is not whether the methods agree exactly, but whether they recover the same qualitative pattern—broad emotional diversity in text alongside visual concentration in calm and happiness.

We note that these alternative approaches differ from our primary LMM approaches in their unit of analysis, so perfect agreement is neither expected nor desired. The CNN image classifier predicts emotions for each individual detected face, whereas the LMM analyzes full images and the characters within them. Similarly, the alternative text measures count emotion-related words (lexicon approach) or classify emotions at the sentence level (transformer method), while our primary text LLM produces character-level emotion labels. What matters is not whether the methods agree exactly, but whether they recover the same qualitative pattern: broad emotional diversity in text alongside visual concentration in calm and happiness. Moreover, the use of different models with varying complexity and transparency provides a source of validation for the findings we obtain leveraging generative AI.

Figure A.6a reports results from the lexicon-based approach. The approach recovers levels of anger, disgust, fear, and sadness comparable to our primary model. The main divergence concerns the calm-happiness split: the LLM identifies roughly 10 percentage points more happiness and 10 percentage points less calm than the lexicon. This may be an artifact of the lexicon approach's sensitivity to keyword choice and the semantic proximity of the two emotions. Indeed, calm and happiness may be lexically difficult to differentiate. For example, consider the word "contentment." It appears in our "calm" dictionary, yet depending on the context of the surrounding text, it could signal either calm or happiness. We thus expect there to be some fluidity between these two categories with the lexicon model. Critically, the aggregate distinction that drives our results – emotional breadth in text – remains intact.

Figure A.6b shows the results of the tranformer-based model. The distribution closely tracks

our primary text results (Figure 2a) with one exception: the GoEmotions model assigns 20% more of its classifications to surprise. This may reflect the difference in the unit of analysis. Operating at the sentence level, the GoEmotions model may, for example, classify sentences associated with discovery and learning as expressing surprise even when no specific character experiences that emotion. (To see this, consider a history textbook chapter on the Age of Discovery.) This expression would not be counted by the LMM because there is no specific character who is exhibiting surprise. Nonetheless, the results from the transformer method are arguably more comparable to those of the LLM than the lexicon-based results since the transformer model and the LLM both incorporate surrounding context to classify emotions rather than relying on isolated keywords.

Figure A.6c shows the results from our alternative image classifier. Consistent with the text robustness checks, the CNN predicts somewhat less calm and more happiness than the LMM for images, but the sum of these two emotions is stable across models. This again reflects the difficulty of distinguishing calm from happiness, now in the visual domain, and reinforces our central image finding: whichever model is used, the two positive emotions jointly dominate visual representation, crowding out the full emotional range. Overall, we find a qualitatively similar distribution to the one obtained using the LMM, which lends further confidence to our main results.

In summary, across classifiers ranging widely in complexity and interpretability, we recover the same qualitative conclusions we obtain from generative AI: text exposes children to a broad emotional repertoire, while images concentrate overwhelmingly on happiness and calm. The specific magnitudes shift with the model and unit of analysis, and calm and happiness prove consistently difficult to separate, both textually and visually, but this ambiguity does not affect our central finding, which concerns the contrast between positive and negative emotions rather than the split within positive emotions. Moreover, because our robustness methods permit alternative units of analysis (word-, sentence-, and face-level), we can confirm that the text-image mismatch is not an artifact of our character-level framing. The convergence of results across independent methods provides meaningful validation of measurements derived from our large generative models.

V.B Manual Validation

In this section, we describe the validation procedures for our primary LMM-based emotion measures. Traditional metrics such as F1 scores or correlations between labels produced by models and humans exhibit several shortcomings. First, these measures do not admit an interpretable threshold above which the models are deemed acceptable. Second, some tasks will give rise to more disagreement between human annotators, and this uncertainty in the ground truth should be incorporated in an assessment of model performance. Third, even in settings where multiple enumerators label the same example, traditional comparisons of the correlations between model and human responses versus those between human and human responses are limited in that comparable correlations may be neither necessary nor sufficient for assessing a model's effectiveness for producing measurements used in downstream applications. Moreover, to the extent that human disagreement is the result of differing but legitimate interpretations of a given example, it may be admissible for

a model to agree with *some* humans, even if it does not agree with all or even the modal annotator.

To address these concerns, we proceed in the spirit of Calderon, Reichart and Dror (2025) by examining how our models compare with human labelers in terms of their ability to agree with other human labelers. Our procedure directly tests the comparability of the model with each human annotator, allowing us to determine whether the model is statistically interchangeable with that annotator. We obtain a "pass rate" (value between zero and one), indicating the fraction of human annotators with whom the model is statistically interchangeable. The higher the pass rate, the more statistically indistinguishable the LMM's classification is from the human annotators' labels. The procedure is outlined in Figure A.7.[8]

More concretely, suppose we have a set of examples denoted $\mathcal{X}$ and human annotators $\mathcal{H}$. For a given example $x \in \mathcal{X}$, let $\mathcal{H}(x)$ denote the set of human annotators who labeled example x. Similarly, for $h \in \mathcal{H}$, let $\mathcal{X}(h)$ denote the set of examples labeled by annotator h. For annotator $h \in \mathcal{H}(x)$, we can also define $\mathcal{H}_{-h}(x)$ to be the set of annotators who labeled example x, excluding annotator h. For $h \in \mathcal{H}(x)$, let $h(x)$ denote the label that human annotator h assigns to example x. As in Calderon, Reichart and Dror (2025), let $f(x)$ denote the model's label assigned to example x. Then, for a given annotator $h_1 \in \mathcal{H}$ who labeled example $x \in \mathcal{X}$, we define the leave-one-out agreement score with annotator $h_2 \in \mathcal{H}$ excluded as

$$\mathrm{ACC}_{h_1}(x, h_2) = \frac{1}{\left|\mathcal{H}_{-h_2}(x)\right|} \sum_{h' \in \mathcal{H}_{-h_2}(x)} \mathbf{1}\{h_1(x) = h'(x)\}.$$

With this definition, we construct a measure of how the model performs compared to human h on example x in terms of relative agreement with all other human annotators who labeled x:

$$d(x, h) = \mathrm{ACC}(f, x, h) - \mathrm{ACC}(h, x, h).$$

Then, we compute

$$\mu(h) = \frac{1}{|\mathcal{X}(h)|} \sum_{x \in \mathcal{X}(h)} d(x, h).$$

We then perform a one-sided t-test of the null $H_0 : \mu(h) \geq 0$. A low (negative) value of $\mu(h)$ provides evidence that the model agrees less with other annotators than annotator h on the set of examples labeled by h.

We can perform the same test for every human annotator. Of course, there exists dependence across the test statistics. We apply the Benjamini-Yekutieli procedure to obtain adjusted p-values that control the false discovery rate under arbitrary dependence among the test statistics. For a fixed level α, we determine the fraction W of adjusted p-values that are above α. This fraction represents the "pass rate," which intuitively is the fraction of human annotators with whom we cannot reject that the model is interchangeable.

[8] We also report traditional macro F1 scores in Table B.3.

We note that our procedure deviates from the alternative annotator test proposed by Calderon, Reichart and Dror (2025) in two important ways. First, they define an example-level win rate in place of our relative agreement. That is, they define

$$\tilde{d}(x,h) = \mathbf{1}\{\text{ACC}(f,x,h) - \text{ACC}(h,x,h) \geq 0\}.$$

This binarization discards information in the sense that the magnitude of the discrepancy between the model's agreement and h's agreement with other humans is not accounted for. The second departure we make from Calderon, Reichart and Dror (2025) is that we test for whether the model exhibits significantly lower agreement with other annotators rather than significantly higher agreement. Due to the fact that having equal agreement with humans is arguably the theoretical limit on how well a model can perform in a classification task, testing whether the model agrees with other human annotators significantly more than humans agree with each other sets an implausibly high bar. As a result, these authors introduce a cost parameter ε that biases the test in favor of the model. Given our definitions, the analog of their implementation would be to test $H_0 : d(x,h) \leq -\varepsilon$. The spirit of this cost parameter is to tolerate lower model performance since the model offers a cheaper source of measurement than human annotators. In practice, however, the choice of ε is somewhat arbitrary. We find it more natural and interpretable to test whether the model is statistically underperforming relative to human annotators and avoid the introduction of this cost parameter.

We display the results of this validation exercise in Table B.4. For each exercise, we show the adjusted p-values and pass rates at different significance levels. For several exercises, the LMM exceeds the 0.5 pass rate proposed by Calderon, Reichart and Dror (2025) as the threshold above which the model can justifiably replace human annotators. We comment that the model generally performs well at measuring demographic characteristics. In many settings (and especially in text), demographic characteristics may simply be unavailable, so we are only assessing model performance on a selected set of examples. Nonetheless, in cases where demographic information can be ascertained, the model performs well.

At the same time, while the model performs comparably to some annotators in terms of emotion labeling, the 0.5 threshold is not always exceeded. The model's emotion labels in text are not statistically distinguishable from two out of six annotators. It performs better in image labeling, meeting or exceeding the 0.5 pass rate at each significance level. To the extent that humans can hold distinct but reasonable interpretations of an emotion, we view these pass rates favorably. However, after binning the emotions into two valence categories ($\{\text{happy, calm}\}$ and $\{\text{any other emotion}\}$, as well as $\{\text{no emotion, unsure}\}$), we see that the pass rate in image labeling declines. While the models may perform satisfactorily overall, these results motivate an analysis to verify that the types of mistakes the models make are not driving our primary results.

To do this, we use the validation sample to estimate the simulated distribution of human

emotion labels we predict would be obtained if humans had labeled the entire dataset. We let $\mathcal{C}$ denote the set of emotion classifications. We observe that

$$\mathbb{P}(h(x) = c) = \sum_{c' \in \mathcal{C}} \mathbb{P}(h(x) = c | f(x) = c') \mathbb{P}(f(x) = c').$$

We obtain $\mathbb{P}(f(x) = c')$ from the model labels on the entire dataset, and we estimate $\mathbb{P}(h(x) = c | f(x) = c')$ using the validation sample. We obtain consistent estimates since the validation sample was obtained by stratified random sampling from the full sample (with LMM class labels as strata to ensure representation of each class). The results are shown in Figure A.8 with 95% confidence intervals obtained via stratified bootstrap resampling from the validation data. We observe that people ascribe calm to more examples in text, and, compared to the model, they label more emotions in faces as happy rather than calm. Nonetheless, we recover the same pattern from the predicted human labels that text displays both positively and negatively valenced emotions while images show happy and calm almost uniformly. This is further confirmed in Figure A.9, which shows the distribution of emotions predicted by humans and the model after binning into Happy/Calm and Non-Happy/Calm.

In Figure A.10, we also show an additional robustness check for the LMM classifications in the text where we subset to emotion labels that the model has high confidence in (≥ 0.8 on a scale of 0 to 1).[9] This is a different subset of characters (likely characters where the emotion is clear). When we examine the overall representation of emotions in this subset, we see the largest changes in the proportion of happiness (10% increase) and confusion (9% decrease), but our main takeaways remain the same.

Overall, the results of the validation exercise lend confidence in the quality of the measurements produced by our LMMs. The model is never statistically worse than all annotators, and it exceeds the prescribed 0.5 threshold in several instances. To the extent that human disagreement reflects legitimate alternative interpretations, this means that the models are not out of line with the labels humans might propose. The model is not identical to human annotators, but we nevertheless recover the same qualitative conclusions after simulating the distribution of emotion labels that we predict humans would produce if they were to label the entire dataset.

VI Consumer Demand and Emotions

Our analysis was motivated by the fact that it is not ex ante obvious how children's materials will represent emotions. Because publishers operate in markets, the emotional content they supply may reflect the types of books purchasers are more likely to select. If purchasers value materials that build emotional intelligence by exposing children to the richness and complexity of real-life

[9] We note that the self-reported confidence score has no probabilistic interpretation, but model performance increases with this score. In particular, at this threshold, we obtain a pass rate of 5/6. We retain 282 validation examples after filtering on the confidence score, so this is not driven by a decrease in power of the test. We also see that the macro F1 score and LLM-human agreement relative to human-human agreement both increase with the confidence threshold.

emotional experience, then publishers would have incentives to produce books that portray a broad emotional repertoire across both text and images. If, instead, purchasers value materials that conform to prevailing social and labor market norms about which emotions are appropriate to display, then publishers would have incentives to emphasize calmness and happiness, especially in visible features such as illustrations and covers. A third possibility is that purchasers may prefer emotional complexity in principle but make choices based on salient, easily processed cues that do not fully capture those preferences. In each case, purchaser behavior can shape the emotional representations children are likely to encounter.

Having documented a systematic oversupply of happy and calm emotions in images relative to the text in two influential sources of children's media, we next ask whether these patterns are consistent with purchaser demand. If purchasers are more likely to select books whose visible features convey positive emotions, then authors, illustrators, and publishers may be incentivized to emphasize positive expressions, especially in images and on covers. This is consistent with other evidence that purchaser demand shapes media content (Gentzkow, Glaeser and Goldin, 2006; Gentzkow and Shapiro, 2010; Sacerdote, Sehgal and Cook, 2020) and that salient product attributes influence choice when purchasers have limited attention (Ellison, 2006; Bordalo, Gennaioli and Shleifer, 2013).

Book covers provide a natural setting to study this mechanism. Covers are highly visible at the point of purchase and can function both as attention-grabbing product attributes and as signals of a book's tone and content (Bloch, 1995; Pieters and Wedel, 2004; Underwood and Klein, 2002; Eliaz and Spiegler, 2011). This is especially relevant for online purchases included in our sample, where purchasers often cannot inspect interior pages before buying. If purchasers favor happy or calm images, this could create market incentives to supply books with these positive emotions on the cover, even when the text contains a broader range of emotions.

This discussion yields a simple empirical prediction: conditional on book characteristics, titles whose covers depict only happy or calm characters should be selected more often than titles whose covers include a character displaying another emotion. We test this prediction using two measures of selection in the children's literature sample: household purchases from the Numerator data and holdings in the Seattle Public Library system.

VI.A Empirical Analysis of Purchaser Demand

We restrict the consumer demand exercise to the children's literature collection, whose consumption more directly reflects purchaser choice. Textbook adoption, by constrast, occurs at the school- or district-level and is largely insulated from individual purchaser choice.[10]

We study purchaser demand using two separate data sources. First, we estimate a conditional logit model using micro-level book purchase data from the Numerator OmniPanel, which records individual buying decisions and thus speaks directly to purchasers' revealed choices. A limitation

[10]Parents can select into schools, but textbook content is only one dimension of this multifaceted decision.

of these data is that we cannot observe purchasers' consideration sets, the books an individual actually considered before buying. Online purchasers (the majority of our sample) could plausibly select any title in our collection, so we address this concern with a second measure: using data on book inventory in the Seattle Public Library system. We expect that a public library would consider purchasing a much broader range of books than a typical individual purchaser, making the assumption that every award-winning book in our sample belongs to the library's consideration set more credible. Moreover, while public libraries may have different objective functions than individual purchasers, we expect that they would generally stock more copies of frequently borrowed books, so library holdings may indirectly reflect reader demand.

To study the disproportionate representation of happy and calm in images specifically, we collapse emotions into two valence categories: happy/calm or non-happy/calm, with the latter category indicates that a character displays some emotion besides happiness or calm.[11] Using the Numerator data, we estimate a conditional logit model of purchasers' decisions. We then use the public-library inventory data to regress the number of copies of a book held by the library against the same set of covariates. Both specifications include an indicator for whether the cover depicts only happy or calm characters, an indicator for whether the cover depicts any character at all, and a vector of book controls: indicator variables for the award(s) the book won, copyright year, whether the book is a picture book, whether the book contains no pictured characters, page count, and average words per page.[12] We do not expect purchasers to accurately perceive the emotional composition of a book's interior. Nonetheless, we estimate two additional specifications that control for the proportions of happy/calm emotions in both images and text by (1) directly including these proportions as linear controls and (2) by adding indicators for quintiles of these proportions within their respective distributions. These controls allow for purchasers who attend more closely to the interior emotional content and may additionally proxy for unobserved confounders.

Table 1 reports the household purchase results from the Numerator data. The emotional representation on the cover plays a substantive role in shaping purchasers' purchase decisions. In particular, across all specifications, purchasers place a significant premium on books depicting happy or calm characters on the cover relative to books depicting other emotions. Notably, purchasers significantly favor books with *no* characters on the cover to books with a character displaying an emotion which is neither happy nor calm. This suggests an explicit aversion to negative emotional imagery: purchasers prefer even an emotionless cover to one signaling anger, fear, or sadness. Such an aversion would, in turn, discourage authors and illustrators from depicting these emotions, generating precisely the visual pattern we document.

In Table 2, we present the analogous results using the Seattle Public Library system data. The findings are qualitatively similar: the library stocks more copies of books with exclusively happy or calm cover characters, or with no character at all, than of books with non-happy/calm

[11] "No emotion" and "unsure" labels are not included in this category.

[12] These last two features serve as a proxy for the age of the target audience of the book.

characters on the cover. Because the library's consideration set is far broader than any household's, this suggests that our inability to observe individual purchasers' consideration sets does not drive Table 1 results. Instead, these results jointly speak to a consistent patterns of readers favoring books with depictions of happiness and calm over books with other emotional depictions.

Additionally, in order to connect our findings to purchasers' decisions about happy and calm emotions in general and not just on covers, we find that the emotions displayed on the cover do provide an informative signal about the emotions in the content of the book. We estimate another conditional logit model to study the probability that a cover contains only happy/calm characters as a function of the proportion of happy/calm emotions in text and images in the content of the book. We also perform an analogous analysis to study the probability a cover contains only non-happy/calm characters. The results are displayed in Table 3. We find that having higher proportions of happy/calm emotions in text and images predicts having exclusively happy/calm characters on the cover. Likewise, lower proportions of happy/calm emotions in text and images predict having exclusively non-happy/calm characters on the cover.[13] These results indicate that covers do provide purchasers with an informative signal about the emotional content within the book (e.g. they can indeed judge a book by its cover). Therefore, purchasers' choosing books with more happy/calm characters pictured on the cover is also consistent with wanting these emotions displayed in the content of the book.

In light of these results, the patterns we uncover in Section IV are consistent with a supply-side response to purchasers' revealed choices regarding emotional displays in children's materials. The evidence does not allow us to distinguish whether purchasers genuinely prefer books that visually emphasize happy and calm emotions, whether they believe such displays are more appropriate for children, or whether they are imperfectly attentive to the broader emotional content of the book. However, all three interpretations help connect the demand results to our motivating empirical question. Rather than rewarding visual representations that expose children to the full range of emotional experience, purchaser behavior appears to reward books whose visible features align with socially valued emotional displays. Moreover, since images – especially those on covers – are particularly salient compared to text, authors would face a stronger disincentive to depict non-happy/calm emotions in the images compared to text. This is consistent with our finding that emotions are relatively diverse in text while images predominantly show characters who are either happy or calm. We also note that this market equilibrium is not necessarily innocuous; if purchasers are boundedly rational but are attracted to depictions of positive emotions, then the resulting equilibrium may lead to an oversupply of happy/calm emotions relative to what is best for young readers.

VII Conclusion

This paper asks an upstream question about the formation of emotional skill: what are children taught to recognize and display through the materials they encounter in and out of school?

[13] The reason the latter result is not mechanical is because these two categories are not mutually exclusive. Covers could have a mix of happy/calm and non-happy/calm emotions, or they could display no characters at all.

Using large multimodal models to classify emotions in public-school textbooks and award-winning children's literature, we document a systematic divergence between what children read and what they see. Text exposes children to a broad emotional repertoire including not only happiness and sadness, but also anger, fear, confusion, and surprise. Images do not: calm and happiness together account for roughly 90% of visually depicted emotions, while emotions traditionally coded as "negative" are almost entirely absent, even on pages whose text describes them. This mismatch is not an artifact of aggregation: it holds within pages, within individual matched characters, and across the full distribution of the corpus, where nearly half of all pages show no overlap between the emotions conveyed in text and those conveyed in images.

These patterns prove remarkably stable. They persist across three decades of textbooks and a century of literature, across states occupying opposite ends of the political spectrum, across subjects and grade levels, and across every character identity we examine. Where identity shapes emotional representation, it does so within the boundaries of positive affect: female characters are more likely to be shown as happy while male characters are more likely to be shown as calm, for instance, subtly reinforcing gendered emotional norms rather than widening the range children encounter.

Our findings speak to several literatures. First, they inform the production function of socioemotional skill. Standard models of human-capital formation emphasize early investments, and children's media constitute one such investment; we show that these materials transmit a narrow set of visually modeled emotions even as their text conveys a broad one. In doing so, we flip the claim that emotions matter for economic choice: rather than asking how emotions shape adult decisions, we ask how emotional display norms are themselves produced through childhood exposure. Second, we contribute to cultural economics by shifting attention by identifying emotional norms, or implicit display rules embedded in cultural goods. Third, we expand prior work on emotional representation in children's books, which has relied largely on manual content analysis of text alone, by developing a scalable multimodal framework that reveals a striking disconnect between modalities. Finally, we contribute to the economics of media and product markets by showing that this pattern may be an equilibrium outcome rather than an exogenous feature of content. Both households and public libraries are significantly less likely to select books whose covers depict "negative" emotions, and covers reliably signal interior emotional content. Because covers are salient and easy to process, authors and illustrators face incentives to oversupply happy and calm imagery, especially in visual features, generating precisely the pattern we document.

Two features of this equilibrium warrant emphasis. First, we cannot fully disentangle whether purchasers genuinely prefer positive imagery, believe such displays are more appropriate for children, or are simply drawn to salient smiling faces under limited attention. All three interpretations connect our demand results to the representational patterns we uncover, but they carry different welfare implications. If purchasers are boundedly rational yet attracted to positive expressions, the resulting equilibrium may oversupply happy and calm emotions relative to what best serves young readers.

Second, the mechanism differs across our two corpora: children's literature responds to a market, whereas textbooks are adopted institutionally, suggesting that the visual constraint in curricular materials reflects curricular design or publisher convention rather than consumer demand alone.

These conclusions come with limits. Our analysis documents representation, not downstream effects; we do not measure how the text-image mismatch shapes children's emotional development, and establishing that link is an important direction for future work. Our emotion classifications, though validated against human annotators and alternative models, remain model-generated, and calm and happiness prove genuinely difficult to separate, though this ambiguity does not affect the positive versus negative contrast that drives our results.

More broadly, our findings underscore the role of children's materials in the formation of emotional skills and norms. Emotional intelligence develops through repeated exposure to emotional language, visual cues, and models of expression. We show that influential educational materials offer children a rich emotional vocabulary in text while visually reinforcing a far narrower repertoire centered on happiness and calm. This pattern is consistent with social and market incentives that reward positive emotional displays, but it may also constrain the range of emotions children see modeled during the very years when emotion-recognition skills develop. From an educational standpoint, the results point to the value of more intentional emotional diversity across both text and images so that children encounter not only examples of emotional regulation, but visual models of the full range of emotions that shape everyday life.

References

Adolphs, Ralph, Leonard Mlodinow, and Lisa Feldman Barrett. 2019. "What is an emotion?" Current Biology, 29(20): R1060–R1064.

Adukia, Anjali, Alex Eble, Emileigh Harrison, Hakizumwami Birali Runesha, and Teodora Szasz. 2023. "What we teach about race and gender: Representation in images and text of children's books." The Quarterly Journal of Economics, 138(4): 2225–2285.

Adukia, Anjali, and Emileigh Harrison. 2025. "Separation of church and state curricula? Examining public and religious private school textbooks." National Bureau of Economic Research.

Angelucci, Manuela, Raissa Fábregas, and Antonia Vazquez. 2026. "The Well-Being Effects of Digital Mental Health Care." ROCKWOOL Foundation Berlin (RFBerlin).

Arizpe, Evelyn, and Morag Styles. 2015. Children reading picturebooks: Interpreting visual texts. Routledge.

Ash, Elliott, Ruben Durante, Maria Grebenschikova, and Carlo Schwarz. 2021. "Visual representation and stereotypes in news media."

Austin, Elizabeth J, Donald H Saklofske, and Sarah M Mastoras. 2010. "Emotional intelligence, coping and exam-related stress in Canadian undergraduate students." Australian Journal of Psychology, 62(1): 42–50.

Averill, James R. 1983. "Studies on anger and aggression: Implications for theories of emotion." American Psychologist, 38(11): 1145.

Bazzi, Samuel, Martin Fiszbein, and Mesay Gebresilasse. 2020. "Frontier culture: The roots and persistence of "rugged individualism" in the United States." Econometrica, 88(6): 2329–2368.

Biasi, Barbara, Minseon Park, John D Singleton, and Seth D Zimmerman. 2025. "Identity and Ideology in the School Boardroom." National Bureau of Economic Research.

Bloch, Peter H. 1995. "Seeking the Ideal Form: Product Design and Consumer Response." Journal of Marketing, 59(3): 16–29.

Bollen, Johan, Huina Mao, and Alberto Pepe. 2011. "Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena." Vol. 5, 450–453.

Bonhomme, Stéphane. 2021. "Teams: Heterogeneity, sorting, and complementarity." arXiv preprint arXiv:2102.01802.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and consumer choice." Journal of Political Economy, 121(5): 803–843.

Bowles, Samuel, and Herbert Gintis. 1976. Schooling in capitalist America. Vol. 57, New York: Basic Books.

Brackett, Marc A, Susan E Rivers, Maria R Reyes, and Peter Salovey. 2012. "Enhancing academic performance and social and emotional competence with the RULER feeling words curriculum." Learning and Individual Differences, 22(2): 218–224.

Brady, Shannon M, Marissa Ogren, and Scott P Johnson. 2024. "Effects of conflicting emotional cues on toddlers' emotion perception." British Journal of Developmental Psychology, 42(3): 376–391.

Calderon, Nitay, Roi Reichart, and Rotem Dror. 2025. "The alternative annotator test for llm-as-a-judge: How to statistically justify replacing human annotators with llms." 16051–16081.

Castex, Gonzalo, and Evgenia Kogan Dechter. 2014. "The changing roles of education and ability in wage determination." Journal of Labor Economics, 32(4): 685–710.

Cockcroft, Marlaina. 2018. "Caldecott and Newbery Medal wins bring instant boost to book sales." School Library Journal, 64(2).

Côté, Stéphane, Anett Gyurak, and Robert W Levenson. 2010. "The ability to regulate emotion is associated with greater well-being, income, and socioeconomic status." Emotion, 10(6): 923.

Cowen, Alan S, and Dacher Keltner. 2017. "Self-report captures 27 distinct categories of emotion bridged by continuous gradients." Proceedings of the National Academy of Sciences, 114(38): E7900–E7909.

Cowen, Alan S, and Dacher Keltner. 2020. "What the face displays: Mapping 28 emotions conveyed by naturalistic expression." American Psychologist, 75(3): 349.

Cusumano, Carlo M, Francesco Fabbri, and Ferdinand Pieroth. 2024. "Competing to commit: Markets with rational inattention." American Economic Review, 114(1): 285–306.

Deming, David J. 2017. "The growing importance of social skills in the labor market." The Quarterly Journal of Economics, 132(4): 1593–1640.

Demszky, Dorottya, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. "GoEmotions: A dataset of fine-grained emotions." arXiv preprint arXiv:2005.00547.

Denham, Susanne A. 2019. "Emotional competence during childhood and adolescence." In Handbook of Emotional Development. 493–541. Springer.

Dickerson, Kelli L, and Jodi A Quas. 2021. "Emotional awareness, empathy, and generosity in high-risk youths." Journal of Experimental Child Psychology, 208: 105151.

Doan, Sy, and Julia Kaufman. 2024. "What Role Do States Play in Selecting K–12 Textbooks?" State Education Standard, 24(1).

Doherty, Martin. 2008. Theory of mind: How children understand others' thoughts and feelings. Psychology Press.

Dröge, Cornelia, and Rene Y Darmon. 1987. "Associative positioning strategies through comparative advertising: Attribute versus overall similarity approaches." Journal of Marketing Research, 24(4): 377–388.

Duchenne de Boulogne, Guillaume-Benjamin. 1990. The Mechanism of Human Facial Expression (Original work published 1862). Cambridge University Press.

Duckworth, Angela L, Christopher Peterson, Michael D Matthews, and Dennis R Kelly. 2007. "Grit: perseverance and passion for long-term goals." Journal of Personality and Social Psychology, 92(6): 1087.

Edin, Per-Anders, Peter Fredriksson, Martin Nybom, and Björn Öckert. 2022. "The rising return to noncognitive skill." American Economic Journal: Applied Economics, 14(2): 78–100.

Eisenberg, Nancy, Amanda Cumberland, and Tracy L Spinrad. 1998. "Parental socialization of emotion." Psychological Inquiry, 9(4): 241–273.

Ekman, Paul. 1989. "The argument and evidence about universals in facial expressions." Handbook of social psychophysiology, 143(164): 2.

Ekman, Paul. 2016. "What scientists who study emotion agree about." Perspectives on Psychological Science, 11(1): 31–34.

Ekman, Paul, E Richard Sorenson, and Wallace V Friesen. 1969. "Pan-cultural elements in facial displays of emotion." Science, 164(3875): 86–88.

Eliaz, Kfir, and Ran Spiegler. 2011. "Consideration sets and competitive marketing." The Review of Economic Studies, 78(1): 235–262.

Ellison, Glenn. 2006. "Bounded rationality in industrial organization." Econometric Society Monographs, 42: 142.

Elster, Jon. 1998. "Emotions and economic theory." Journal of Economic Literature, 36(1): 47–74.

Entorf, Antonia K, and Thomas J Dohmen. 2025. "The effect of compulsory education on non-cognitive skills: Evidence from low-and middle-income countries." Economics of Education Review, 107: 102654.

Foulds, Kim. 2023. "Co-viewing mass media to support children and parents' emotional ABCs: An evaluation of Ahlan Simsim." Early Childhood Education Journal, 51(8): 1479–1488.

Genevsky, Alexander, and Brian Knutson. 2015. "Neural affective mechanisms predict market-level microlending." Psychological Science, 26(9): 1411–1422.

Gentzkow, Matthew, and Jesse M Shapiro. 2010. "What drives media slant? Evidence from US daily newspapers." Econometrica, 78(1): 35–71.

Gentzkow, Matthew, Edward L Glaeser, and Claudia Goldin. 2006. "The rise of the fourth estate. How newspapers became informative and why it mattered." In Corruption and Reform: Lessons from America's Economic History. 187–230. University of Chicago Press.

Green, Clarence, and He Sun. 2025. "Picturebooks increase the frequency and diversity of emotion vocabulary in children's language environments: Modeling potential benefits to emotional literacy, with pedagogical resources." Early Education and Development, 36(3): 568–586.

Gross, James J. 2015. "Emotion regulation: Current status and future prospects." Psychological Inquiry, 26(1): 1–26.

Gross, James J, and Oliver P John. 2003. "Individual differences in two emotion regulation processes: implications for affect, relationships, and well-being." Journal of Personality and Social Psychology, 85(2): 348.

Groves, Melissa Osborne. 2005. "How important is your personality? Labor market returns to personality for women in the US and UK." Journal of Economic Psychology, 26(6): 827–841.

Gurney, Edmund. 1884. "What is an Emotion?" Mind, 9(35): 421–426.

Hansen, Stephen, Tejas Ramdas, Raffaella Sadun, and Joe Fuller. 2021. "The demand for executive skills." National Bureau of Economic Research.

Hasan, Maryam, Elke Rundensteiner, and Emmanuel Agu. 2014. "Emotex: Detecting emotions in twitter messages."

Heckman, James J, and Yona Rubinstein. 2001. "The importance of noncognitive skills: Lessons from the GED testing program." American Economic Review, 91(2): 145–149.

Heckman, James J, Jora Stixrud, and Sergio Urzua. 2006. "The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior." Journal of Labor Economics, 24(3): 411–482.

Hochschild, Arlie Russell. 1983. "The Managed Heart: Commercialization of Human Feeling."

Hoemann, Katie, Fei Xu, and Lisa Feldman Barrett. 2019. "Emotion words, emotion concepts, and emotional development in children: A constructionist hypothesis." Developmental Psychology, 55(9): 1830.

Hoffman, Mitchell, and Steven Tadelis. 2021. "People management skills, employee attrition, and manager rewards: An empirical analysis." Journal of Political Economy, 129(1): 243–285.

Izard, Carroll E. 1977. "Differential emotions theory." In Human Emotions. 43–66. Springer.

James, William. 1884. "What is Emotion?" Readings in the History of Psychology (1948), pp. 290–303. Appleton-Century-Crofts.

Kardum, Rona Bušljeta, Katarina Dadic, and Martina Horvat. 2021. "Education for the cultivation of emotions through textbooks: the example of Croatian high school history textbooks." Educational Process: International Journal, 10(2): 28–41.

Kautz, Tim, James J Heckman, Ron Diris, Bas Ter Weel, and Lex Borghans. 2014. "Fostering and measuring skills: Improving cognitive and non-cognitive skills to promote lifetime success." , (Working Paper 20749).

Kemper, Theodore D. 1978. "Toward a sociology of emotions: Some problems and some solutions." The American Sociologist, 30–41.

Kumschick, Irina R, Luna Beck, Michael Eid, Georg Witte, Gisela Klann-Delius, Isabella Heuser, Rüdiger Steinlein, and Winfried Menninghaus. 2014. "Reading and feeling: The effects of a literature-based intervention designed to increase emotional competence in second and third graders." Frontiers in Psychology, 5: 1448.

Ku, Yunjung. 2019. "Collocations of English Emotion Expressions in High School Textbooks." Modern English Education, 20(4): 63–73.

Loewenstein, George. 2000. "Emotions in economic theory and economic behavior." American Economic Review, 90(2): 426–432.

Lucas, Richard E, Ed Diener, and Randy J Larsen. 2003. "Measuring positive emotions." In Positive Psychological Assessment: A Handbook of Models and Measures. , ed. Shane J Lopez and Charles Richard Snyder. American Psychological Association.

Ludwig, Jens, and Sendhil Mullainathan. 2024. "Machine learning as a tool for hypothesis generation." The Quarterly Journal of Economics, 139(2): 751–827.

Lundberg, Shelly. 2017. "Non-cognitive skills as human capital." Education, Skills, and Technical Change: Implications for Future US GDP Growth, 77.

Lutz, Catherine, and Geoffrey M White. 1986. "The anthropology of emotions." Annual Review of Anthropology, 405–436.

Marcus, George E. 2000. "Emotions in politics." Annual Review of Political Science, 3(1): 221–250.

Mayer, John D, and Glenn Geher. 1996. "Emotional intelligence and the identification of emotion." Intelligence, 22(2): 89–113.

Michalopoulos, Stelios, and Christopher Rauh. 2024. "Movies." , (Working Paper 32220). Accessed: 2024-08-06.

Michalopoulos, Stelios, and Melanie Meng Xue. 2021. "Folklore." The Quarterly Journal of Economics, 136(4): 1993–2046.

Nikolajeva, Maria. 2014. "33 Emotion ekphrasis: representation of emotions in children's picturebooks." Visual Communication, 711.

Nikolajeva, Maria. 2017. "Emotions in picturebooks." In The Routledge Companion to Picturebooks. 110–118. Routledge.

Öhman, Arne. 2007. "Has evolution primed humans to "beware the beast"?" Proceedings of the National Academy of Sciences, 104(42): 16396–16397.

Öhman, Arne. 2009. "Of snakes and faces: An evolutionary perspective on the psychology of fear." Scandinavian Journal of Psychology, 50(6): 543–552.

Park, BoKyung, Alexander Genevsky, Brian Knutson, and Jeanne Tsai. 2020. "Culturally valued facial expressions enhance loan request success." Emotion, 20(7): 1137.

Parker, James DA, Laura J Summerfeldt, Marjorie J Hogan, and Sarah A Majeski. 2004. "Emotional intelligence and academic success: Examining the transition from high school to university." Personality and Individual Differences, 36(1): 163–172.

Perry, Nicole B, Jessica M Dollar, Susan D Calkins, Susan P Keane, and Lilly Shanahan. 2020. "Maternal socialization of child emotion and adolescent adjustment: Indirect effects through emotion regulation." Developmental Psychology, 56(3): 541.

Pieters, Rik, and Michel Wedel. 2004. "Attention capture and transfer in advertising: Brand, pictorial, and text-size effects." Journal of Marketing, 68(2): 36–50.

Plutchik, Robert, and Henry Kellerman. 1980. Emotion Theory, Research, and Experience. Academic Press.

Rasmussen, Eric E, Autumn Shafer, Malinda J Colwell, Shawna White, Narissra Punyanunt-Carter, Rebecca L Densley, and Holly Wright. 2016. "Relation between active mediation, exposure to Daniel Tiger's Neighborhood, and US preschoolers' social and emotional development." Journal of Children and Media, 10(4): 443–461.

Rorty, Amélie Oksenberg, et al. 1980. Explaining Emotions. Vol. 22, University of California Press Berkeley.

Ruffman, Ted, Qiuyi Kong, Hui Mei Lim, Kangning Du, and Emilia Tiainen. 2023. "Recognition of facial emotions across the lifespan: 8-year-olds resemble older adults." British Journal of Developmental Psychology, 41(2): 128–139.

Russell, James A. 2003. "Core affect and the psychological construction of emotion." Psychological Review, 110(1): 145.

Sacerdote, Bruce, Ranjan Sehgal, and Molly Cook. 2020. "Why is all COVID-19 news bad news?" National Bureau of Economic Research.

Smith, Vicky. 2013. "The 'Caldecott effect'." Children and Libraries: The Journal of the Association for Library Service to Children, 1(1): 9–13.

Solomon, Robert C. 1976. The Passions: The Myth and Nature of Human Emotion. Anchor.

Sullivan, Denis G, and Roger D Masters. 1988. ""Happy warriors"': Leaders' facial displays, viewers' emotions, and political support." American Journal of Political Science, 345–368.

Szasz, Teodora, Emileigh Harrison, Ping-Jung Liu, Ping-Chang Lin, Hakizumwami Birali Runesha, and Anjali Adukia. 2022. "Measuring representation of race, gender, and age in children's books: Face detection and feature classification in illustrated images." Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision.

Thelwall, Mike, Kevan Buckley, and Georgios Paltoglou. 2011. "Sentiment in Twitter events." Journal of the American Society for Information Science and Technology, 62(2): 406–418.

Thoits, Peggy A. 1989. "The sociology of emotions." Annual Review of Sociology, 15(1): 317–342.

Tsai, Jeanne L. 2007. "Ideal affect: Cultural causes and behavioral consequences." Perspectives on Psychological Science, 2(3): 242–259.

Tsai, Jeanne L, Brian Knutson, and Helene H Fung. 2006. "Cultural variation in affect valuation." Journal of Personality and Social Psychology, 90(2): 288.

Underwood, Robert L, and Nicole M Klein. 2002. "Packaging as brand communication: Effects of product pictures on consumer responses to the package and brand." Journal of Marketing Theory and Practice, 10(4): 58–68.

Widen, Sherri C, and James A Russell. 2008. "Children acquire emotion categories gradually." Cognitive Development, 23(2): 291–312.

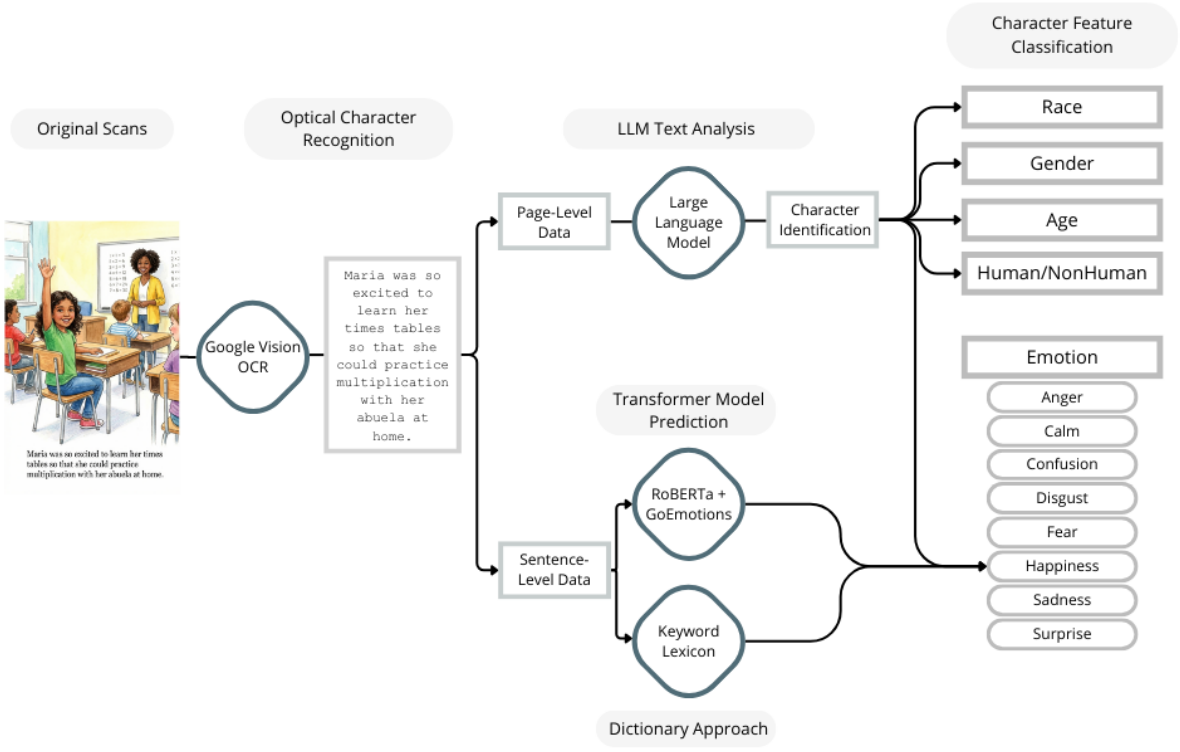
Yanagizawa-Drott, David H., and Hans-Joachim Voth. 2023. "Image(s)." Working Paper, University of Zurich.

Zemach, Eddy M. 2001. "What is emotion?" American Philosophical Quarterly, 38(2): 197–207.

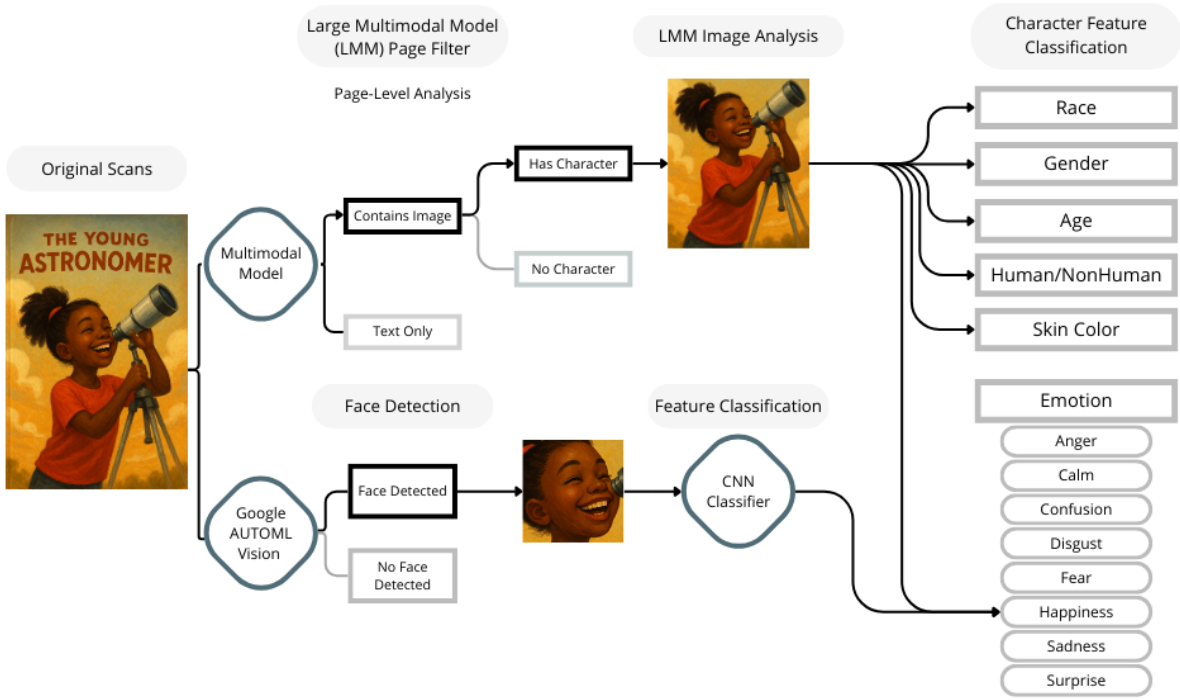
Ziems, Caleb, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. "Can large language models transform computational social science?" Computational Linguistics, 50(1): 237–291.

FIGURE 1
Data Extraction Pipelines

(a) Text Analysis



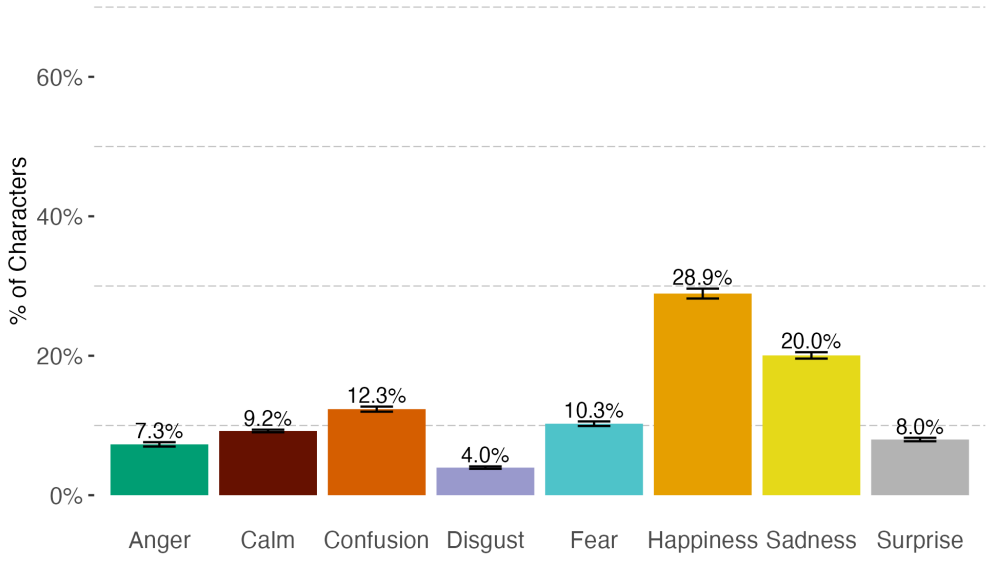
(b) Image Analysis



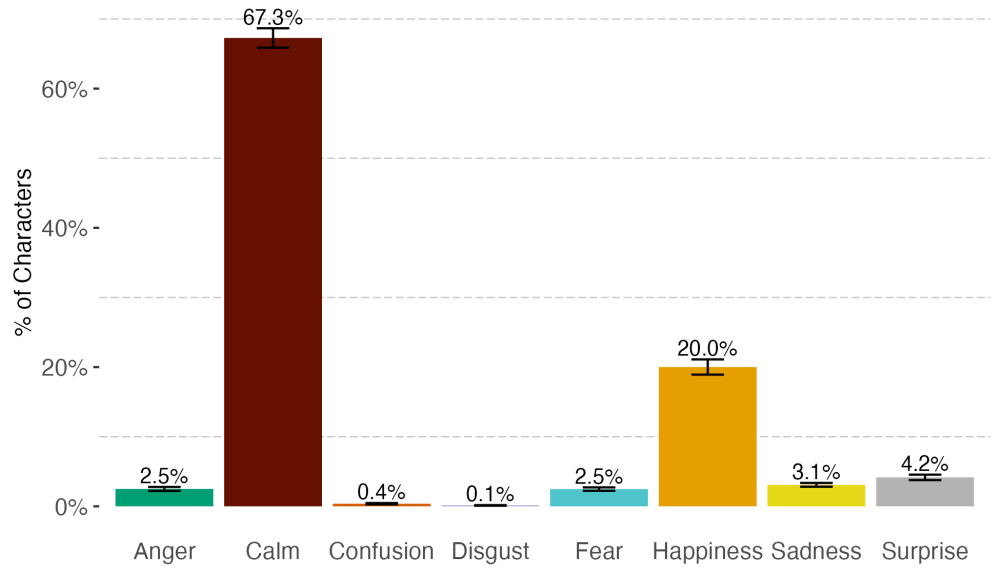
Notes: This figure illustrates the steps in our text analysis (top) and image analysis (bottom) pipelines beginning with the raw scan of each page in our sample of books.

FIGURE 2
Overall Representation of Emotions

(a) Text

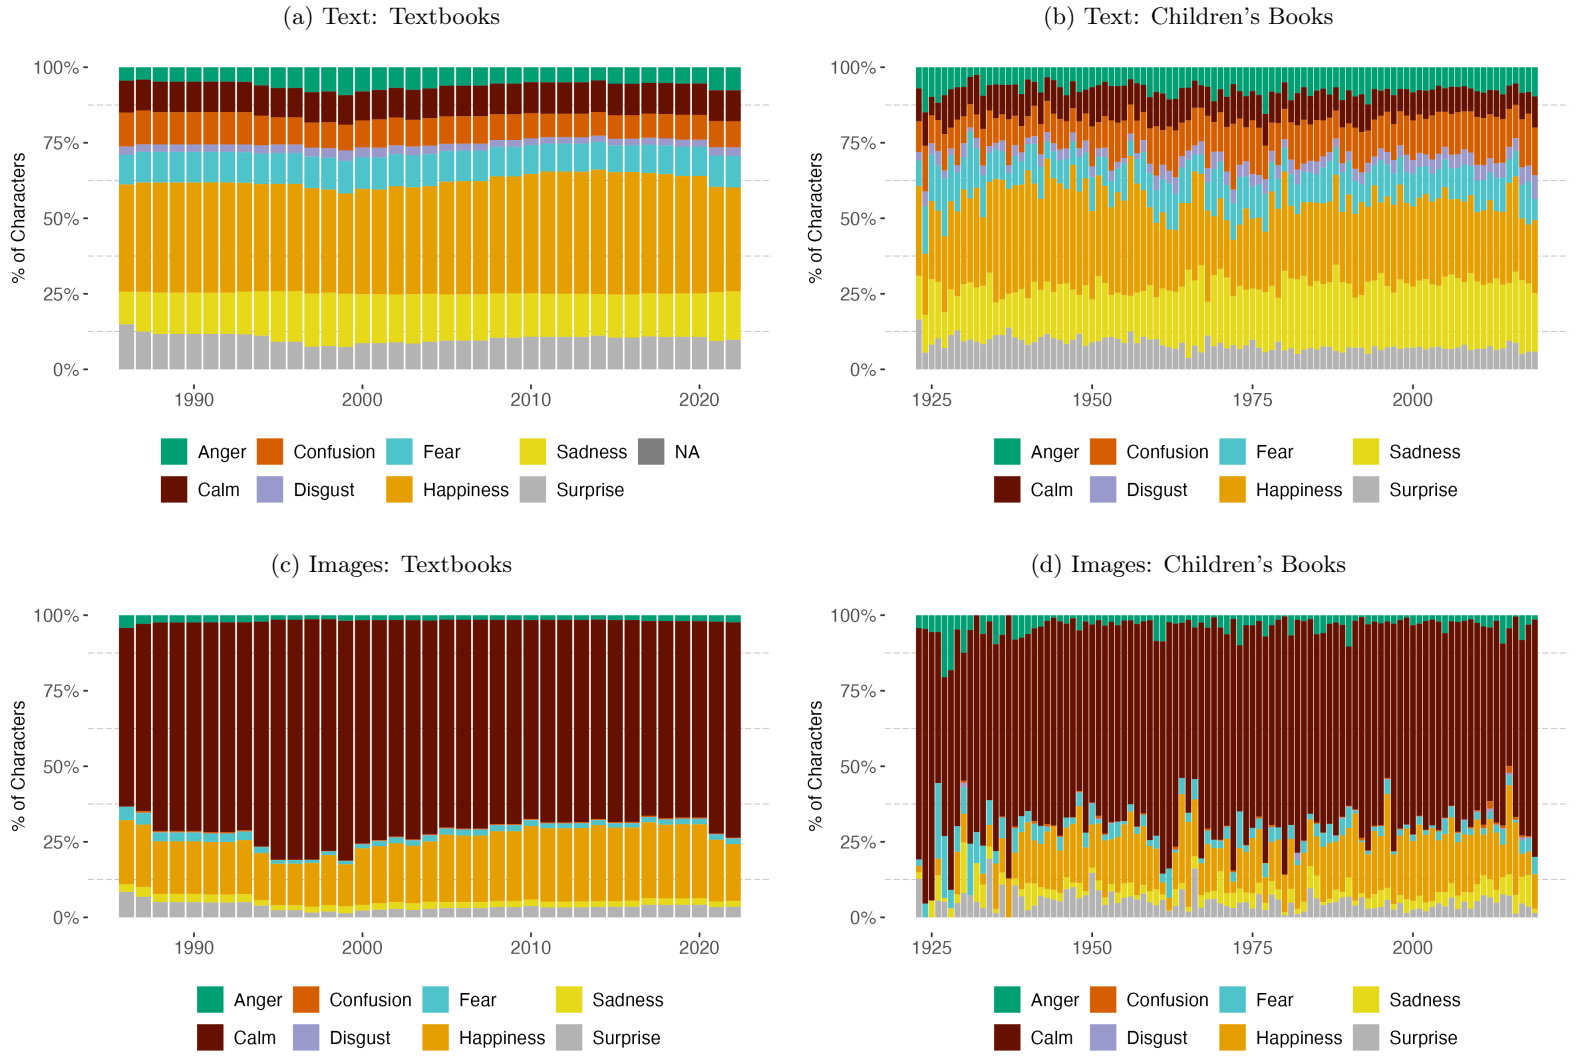


(b) Images



Notes: This figure shows the aggregate displays of emotions in text and images. The top panel shows the percentage of emotions conveyed by characters as a share of all emotions detected in the text of the books (literature and textbooks) using a large language model. The bottom panel shows the percentage of emotions conveyed by faces as a share of all emotions detected in the images of books (literature and textbooks) using a large multimodal model. Error bars depict 95% confidence intervals constructed from standard errors that are clustered at the book level.

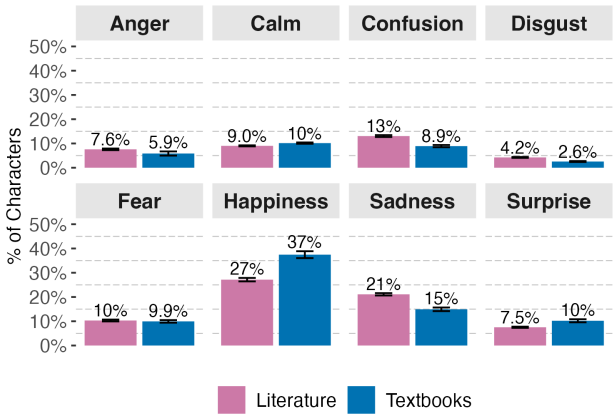
FIGURE 3
Representation of Emotions Over Time



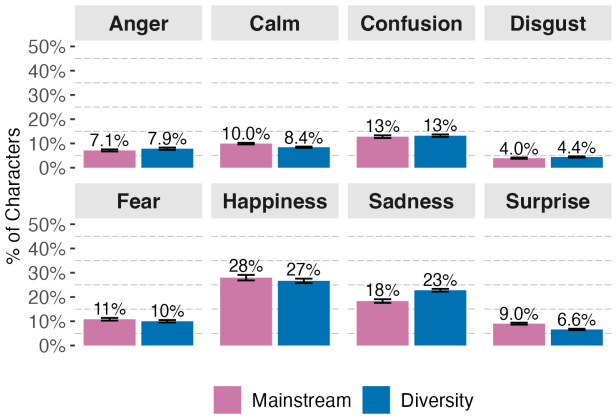
Notes: This figure shows the percentage of emotions detected over time. Panel (a) shows the share of all emotions detected in the sample of textbooks, and panel (b) shows the shares in literature. Emotions are detected using OpenAI. Emotions for each textbook are counted for each year the textbook is in use, and emotions for each children's book are counted for the book's copyright year. Panels (c) and (d) show the analogous results for emotions detected in faces in both corpora. Children's books are associated with their respective copyright years, whereas textbooks are associated with each year they are in use. Note that textbooks are in use for multiple years, whereas the children's books only appear one time in the time series panel.

FIGURE 4
Representation of Emotions in the Text of Different Subcollections

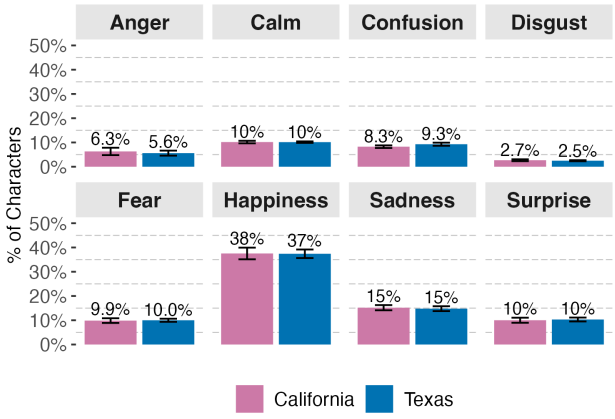
(a) Children's Literature vs. Textbooks



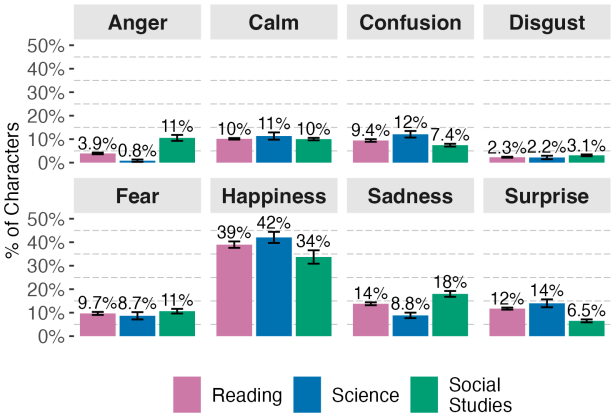
(b) Mainstream vs. Diversity



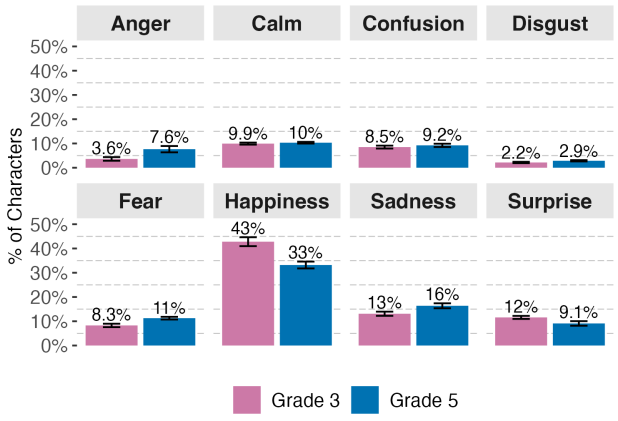
(c) California vs. Texas Textbooks



(d) By Textbook Subject



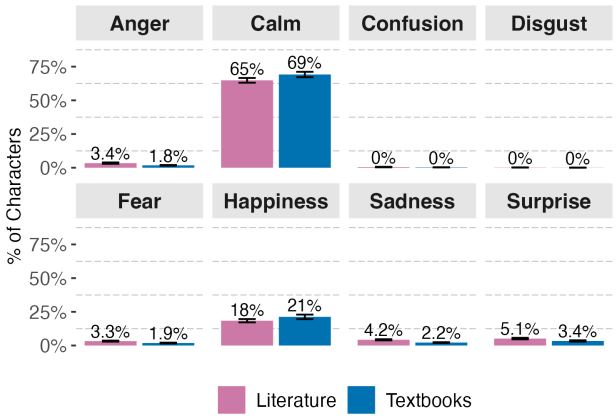
(e) By Textbook Grade



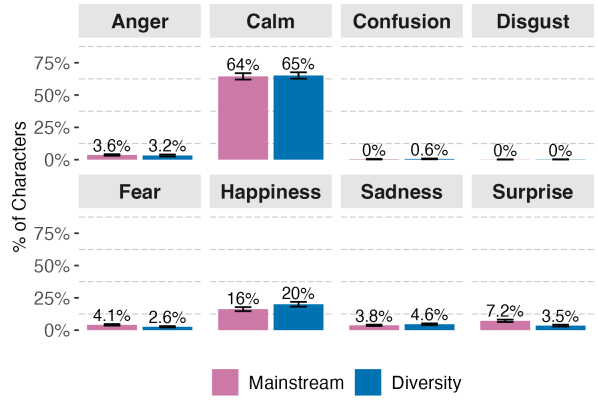
Notes: This figure shows the percentage of emotions conveyed by characters as a share of all emotions detected by OpenAI separately for different subcollections. Panel (a) compares textbooks and children's books, panel (b) compares California and Texas textbooks, panel (c) compares science, reading, and social studies textbooks, panel (d) compares grade three and grade five textbooks, and panel (e) compares the Mainstream and Diversity collections of children's books. Error bars depict 95% confidence intervals constructed from standard errors that are clustered at the book level.

FIGURE 5
Representation of Emotions in the Images of Different Subcollections

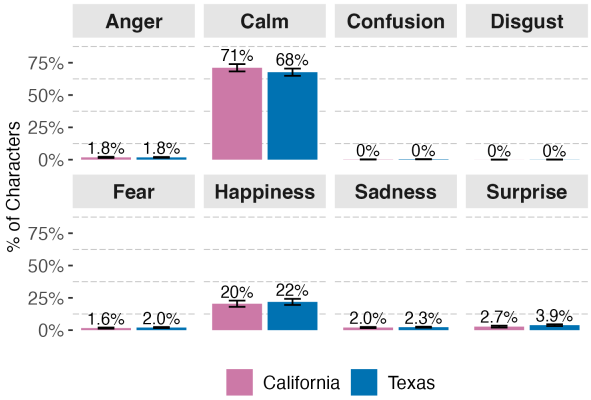
(a) Children's Literature vs. Textbooks



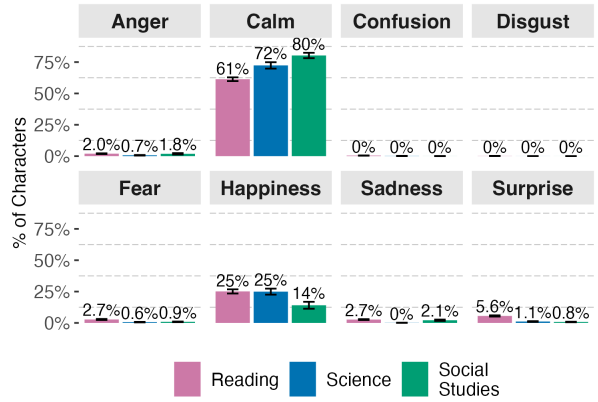
(b) Mainstream vs. Diversity



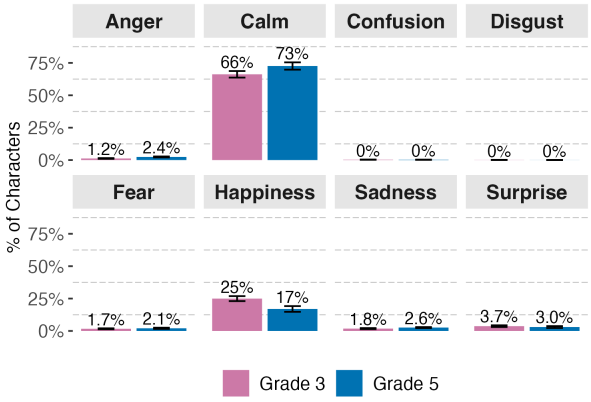
(c) California vs. Texas Textbooks



(d) By Textbook Subject



(e) By Textbook Grade

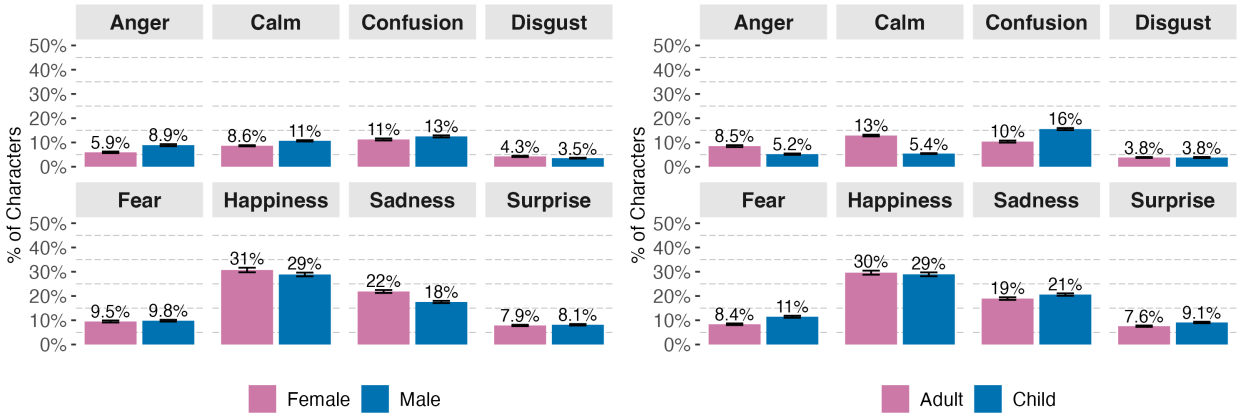


Notes: This figure shows the percentage of emotions conveyed by detected faces as a share of all emotions detected by an LMM separately for different subcollections. Panel (a) compares textbooks and children's books, panel (b) compares California and Texas textbooks, panel (c) compares science, reading, and social studies textbooks, panel (d) compares grade three and grade five textbooks, and panel (e) compares the Mainstream and Diversity collections of children's books. Error bars depict 95% confidence intervals constructed from standard errors that are clustered at the book level.

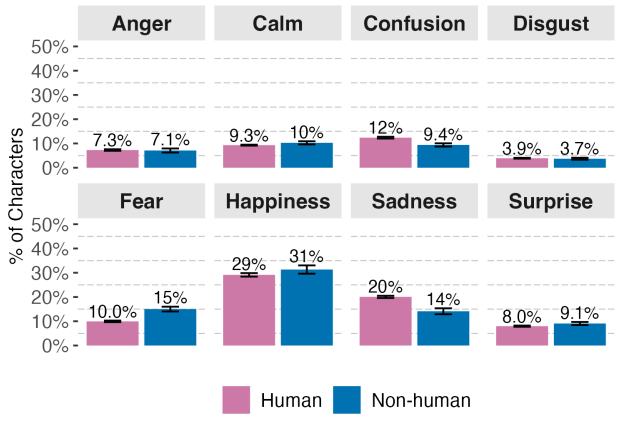
FIGURE 6
Displays of Emotions in the Text for Characters with Different Identities

(a) Gender

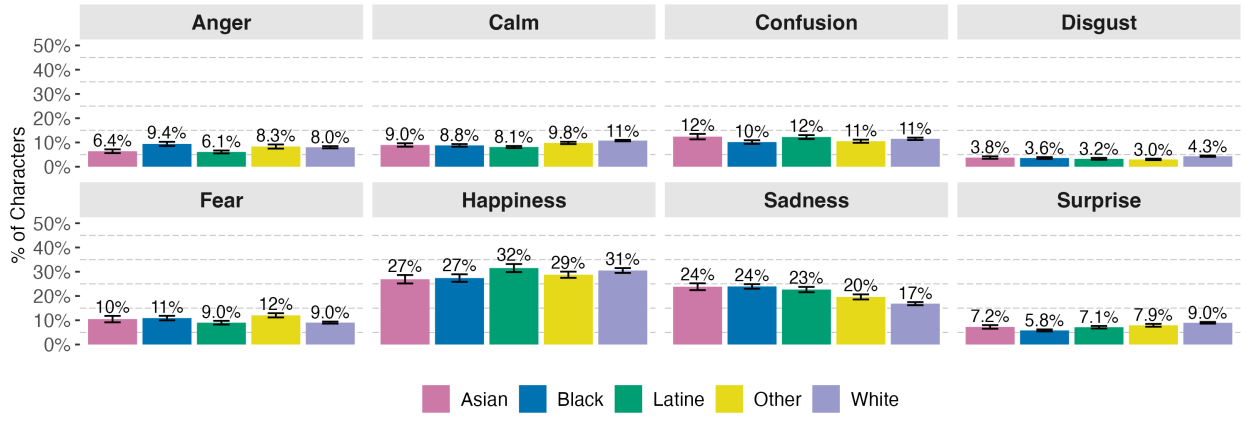
(b) Age



(c) Human vs. Non-Human



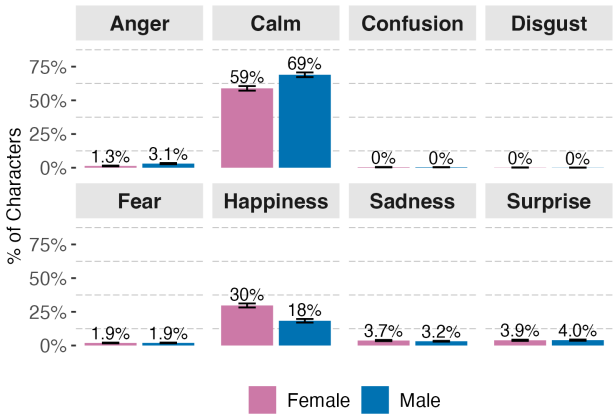
(d) Race



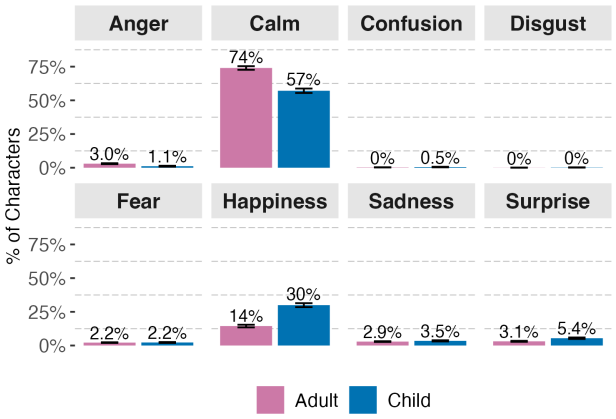
Notes: This figure shows the percentage of emotions conveyed in the text by characters of different identities. Panel (a) compares female and male characters, panel (b) compares human characters of different races, panel (c) compares characters of different ages, and panel (d) compares human and non-human characters. Error bars depict 95% confidence intervals constructed from standard errors that are clustered at the book level.

FIGURE 7
Displays of Emotions in the Faces of Characters with Different Identities

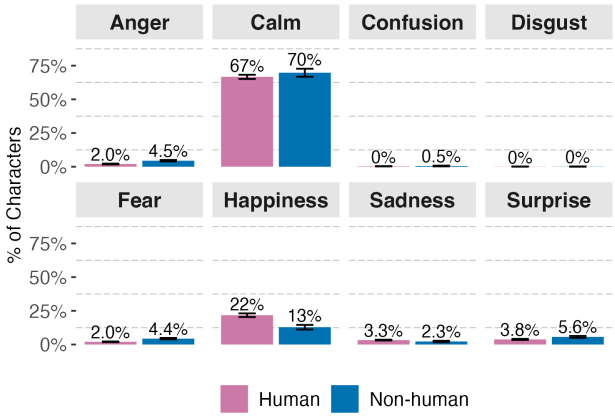
(a) Gender



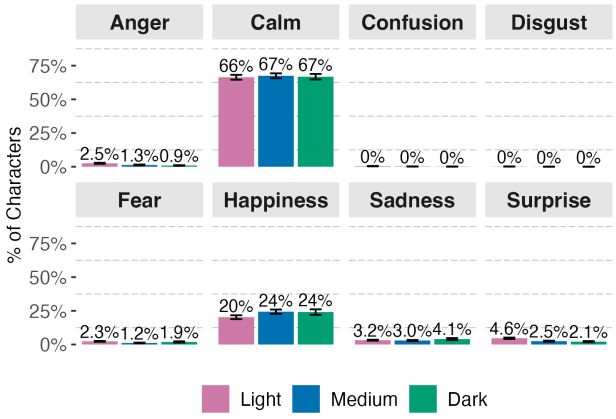
(b) Age



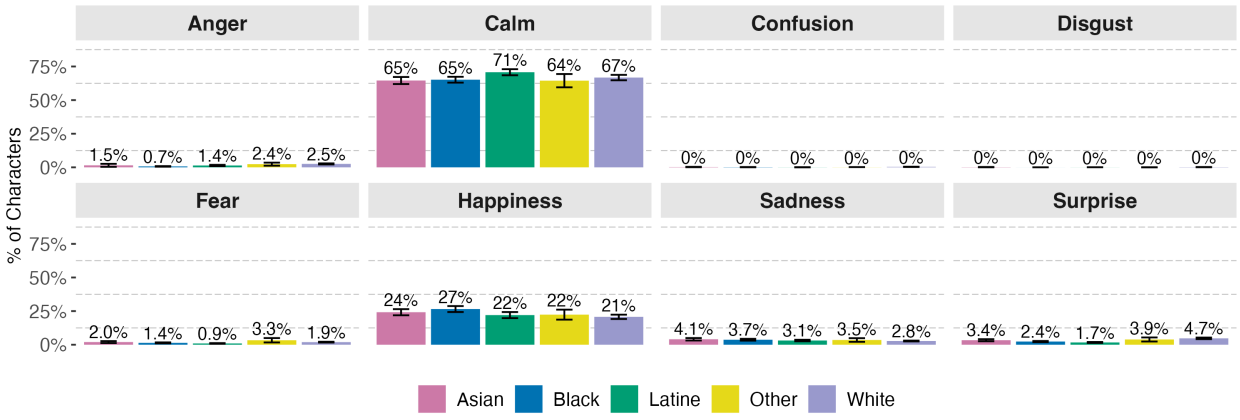
(c) Human vs. Non-Human



(d) Skin Color

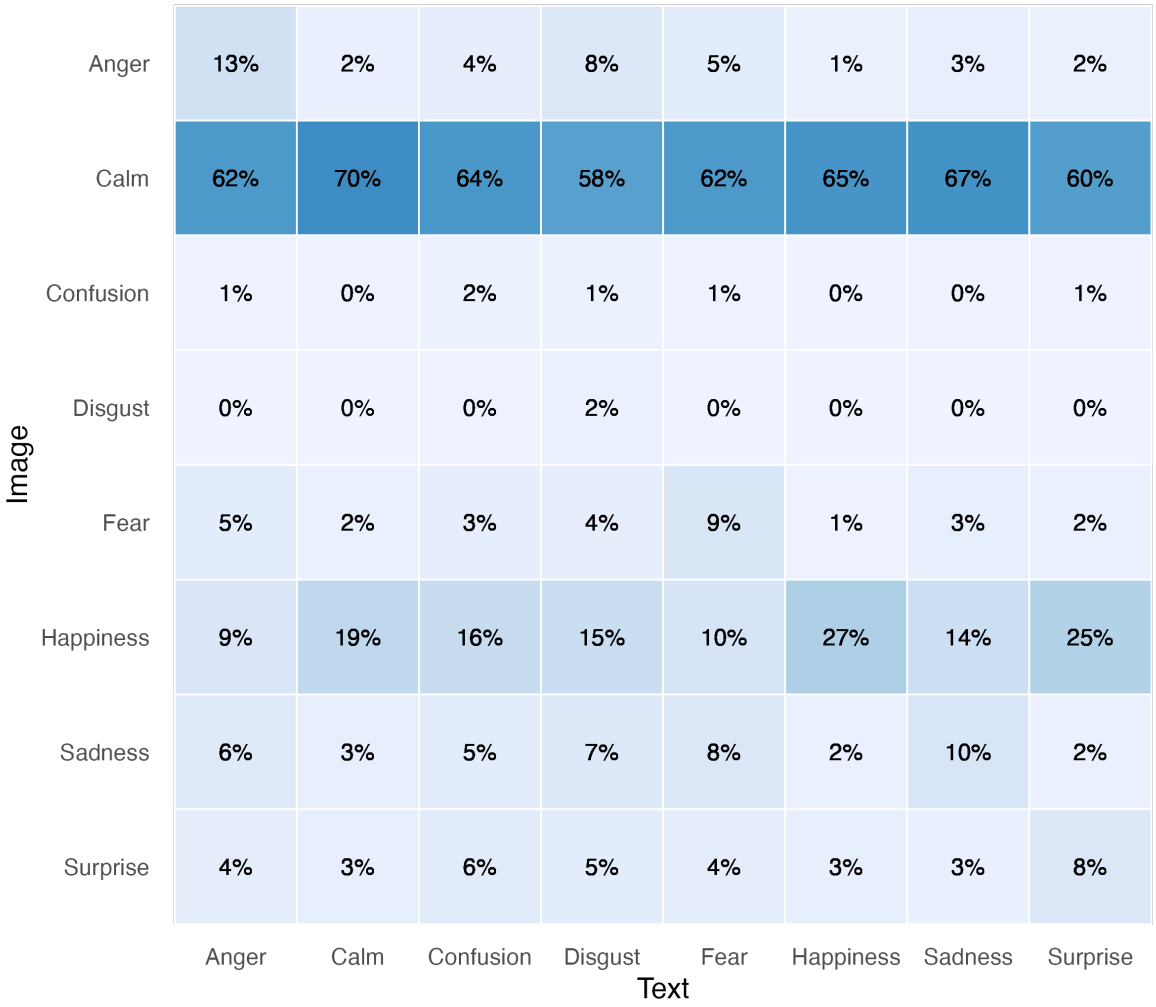


(e) Race



Notes: This figure shows the percentage of emotions conveyed in the detected faces of characters of different identities. Panel (a) compares female and male characters, panel (b) compares human characters of different races, panel (c) compares characters of different ages, panel (d) compares human and non-human characters, and panel (e) compares characters with different skin color. Error bars depict 95% confidence intervals constructed from standard errors that are clustered at the book level.

FIGURE 8
Normalized Co-occurrence Matrix



Note: This figure shows the co-occurrence of emotions for text and images on the same page, normalized by the text columns. In other words, each cell shows the proportion of times a specific image label occurred, given a particular text label. For example, the top left cell shows that when anger is found in text, an image on the same page will depict anger 22% of the time.

TABLE 1
Book Purchases

	Dependent variable: Whether Purchased		
	(1)	(2)	(3)
Happy or Calm Only Cover	0.160***	0.554***	0.667***
	(0.025)	(0.031)	(0.031)
No Cover Character	1.186***	0.490***	0.542***
	(0.024)	(0.030)	(0.031)
Book Controls		X	X
Linear Emotion Controls		X	
Quintile Emotion Controls			X
Observations	28,521	28,521	28,521
AIC	351,716.6	296,521.2	295,216.7
BIC	351,745.1	298,218.3	296,999.3
RMSE	0.05	0.05	0.05

Notes: This table presents results from a logistic regression of the binary purchase outcome on cover emotion measures. The sample is restricted to literature books and characters that express a defined emotion. "Happy or Calm Only Cover" equals 1 when the cover's characters display only calmness or happiness; "No Cover Character" equals 1 when no characters are depicted on the cover; the omitted group is books that contain at least one character who portrays an emotion other than happy or calm. Book controls include a variety of book characteristics including indicator variables for the award(s) the book won, the book's copyright year, the number of pages, average words per page, whether the book is a picture book (characters with emotions only appear in images), and whether it contains no images (characters with emotions appear only in text). Observations refer to the number of times we observe a unique book being purchased by a household across all households. We cluster standard errors by household. In this table, ***, **, and * represent significance at the 0.01, 0.05, and 0.1 levels, respectively.

Researcher(s)' own analyses calculated (or derived) based in part on data from Market Track, LLC dba Numerator and marketing databases provided through the Numerator Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Numerator data are those of the researcher(s) and do not reflect the views of Numerator. Numerator is not responsible for and had no role in analyzing or preparing the results reported herein. LLMs and AI tools are never applied to the Numerator data.

TABLE 2
Library Inventory

	Dependent variable: Number of Books in Branch Inventory		
	(1)	(2)	(3)
Happy or Calm Only Cover	0.1334***	0.1174***	0.0760***
	(0.0191)	(0.0280)	(0.0270)
No Cover Character	0.1213***	0.1839***	0.1366***
	(0.0206)	(0.0420)	(0.0372)
Book Controls		X	X
Linear Emotion Controls		X	
Quintile Emotion Controls			X
Observations	6,944	6,944	6,944
Squared Correlation	0.00313	0.38340	0.38359
Pseudo R^2	0.00132	0.20337	0.20350
BIC	16,514.9	14,196.2	14,247.1

Notes: This table presents results from a linear regression of the library checkout outcome on cover and content emotion measures, controlling for the distribution of emotions in both text and images. The sample is restricted to literature books and characters that express a defined emotion. "Happy or Calm Only Cover" equals 1 when the cover's characters display only calmness or happiness; "No Cover Character" equals 1 when no characters are depicted on the cover; the omitted group is books that contain at least one character who portrays an emotion other than happy or calm. Book controls include a variety of book characteristics including indicator variables for the award(s) the book won, the copyright year of the book, the number of pages, average words per page, whether the book is a picture book (characters with emotions only appear in images), and whether it contains no images (characters with emotions appear only in text). We also add fixed effects for the specific library and cluster standard errors by specific library. In this table, ***, **, and * represent significance at the 0.01, 0.05, and 0.1 levels, respectively.

TABLE 3
Relationship Between the Book Cover and Content

	Book Cover:		
Proportion Happy or Calm in Book Content:	Happy or Calm Cover (1 = yes) (1)	Mixed Cover (1 = yes) (2)	Not Happy or Calm Cover (1 = yes) (3)
Images	3.420***	-3.370***	-5.520***
	(0.471)	(0.867)	(0.982)
Text	2.202***	1.403	-3.510*
	(0.659)	(1.250)	(1.803)
Book Controls	X	X	X
Observations	1,113	1,113	1,113
AIC	1,310.5	557.3	461.4
BIC	1,917.3	1,164.1	1,068.2
Log Likelihood	-534.246	-157.658	-109.700
RMSE	0.40	0.20	0.17

Notes: This table presents results from a logistic regression where the outcome in column one indicates whether a book cover has only happy or calm characters, the outcome in column two indicates whether a book cover has a mix of happy or calm characters and characters that are neither happy nor calm, and the outcome in column three indicates whether a book cover has exclusively characters that are neither happy nor calm. The two variables shown are the proportion of happy or calm characters in images and text in the entire book. The book controls include the number of pages, average words per page, whether there are no characters in the text or images with an emotion, and indicator variables for the award(s) the book won and the copyright year. In this table, ***, **, and * represent significance at the 0.01, 0.05, and 0.1 levels, respectively.

Supplementary Appendix

(How) Do We Teach Emotions?*

July 23, 2026

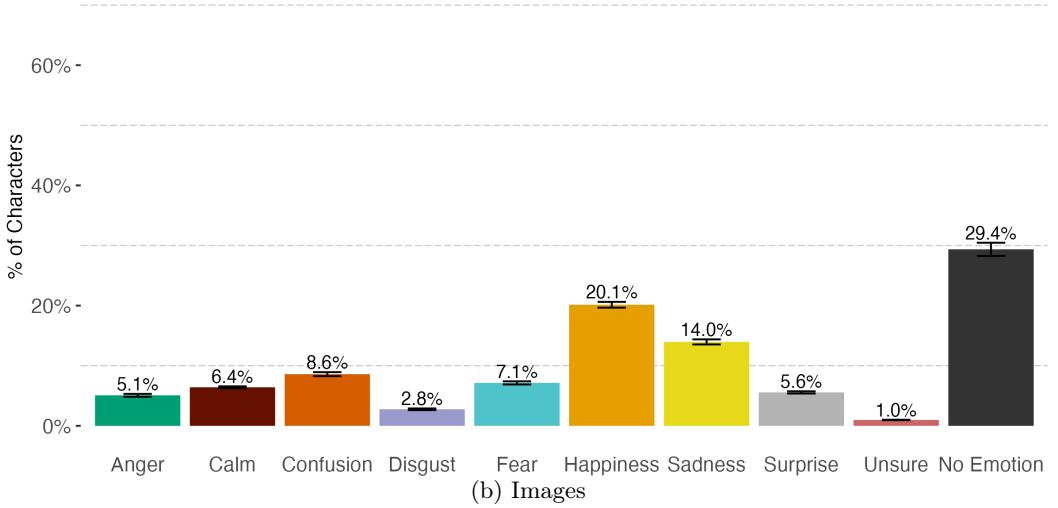
Anjali Adukia
Matthew Bonci
Paula Dastres
Emileigh Harrison
Jake Nicoll
Teodora Szasz

*Contact: adukia@uchicago.edu, mbonci@uchicago.edu, dastresgalla@wisc.edu, emileigh.harrison@austin.utexas.edu, jwnicoll@uchicago.edu, dorha.szasz@gmail.com

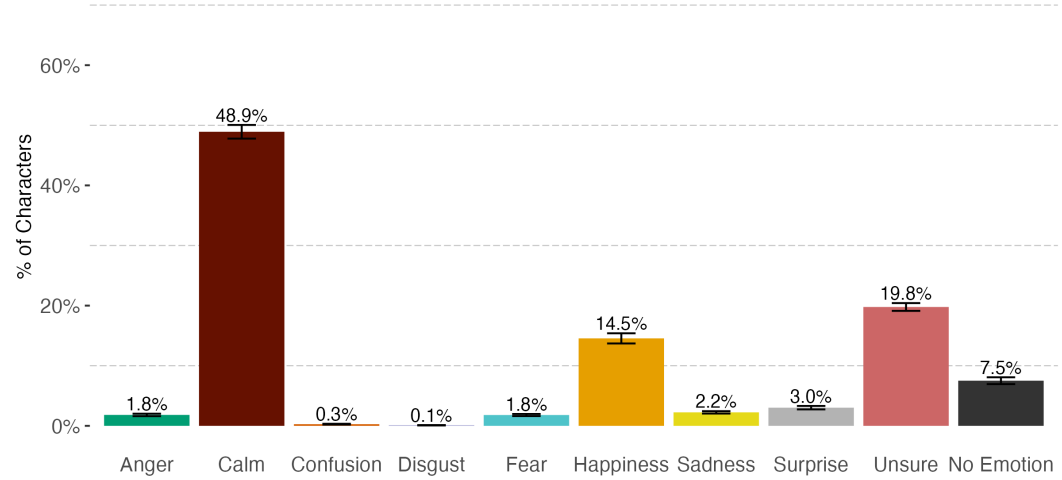
A Appendix Figures

FIGURE A.1
Overall Representation of Emotions (Including No Emotion and Unsure)

(a) Text



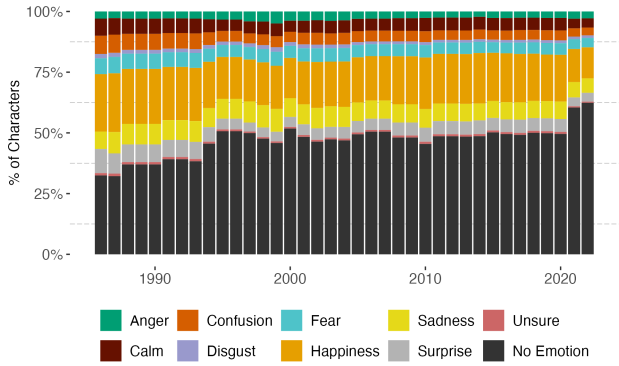
(b) Images



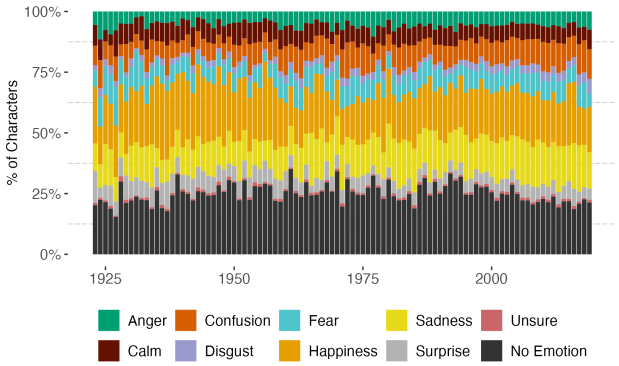
Notes: This figure replicates Figure 2 with additional categories for "No Emotion" and "Unsure." The top panel shows the percentage of emotions conveyed by characters as a share of all emotions detected in the sample of books (literature and textbooks). Emotions are detected using an LMM. The bottom panel shows the percentage of emotions conveyed by faces as a share of all emotions detected in the sample of books (literature and textbooks). Emotions are detected using an LMM.

FIGURE A.2
Representation of Emotions Over Time (including No Emotion and Unsure)

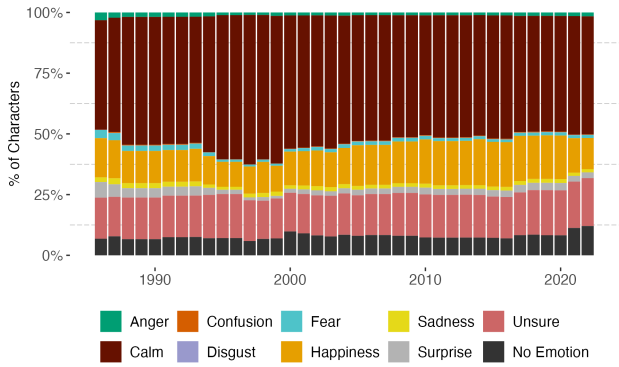
(a) Text: Textbooks



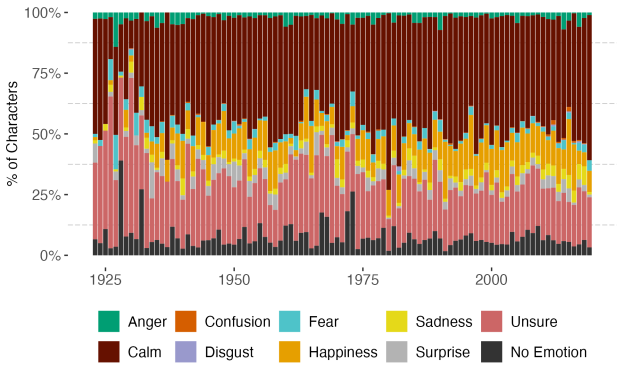
(b) Text: Children's Books



(c) Images: Textbooks



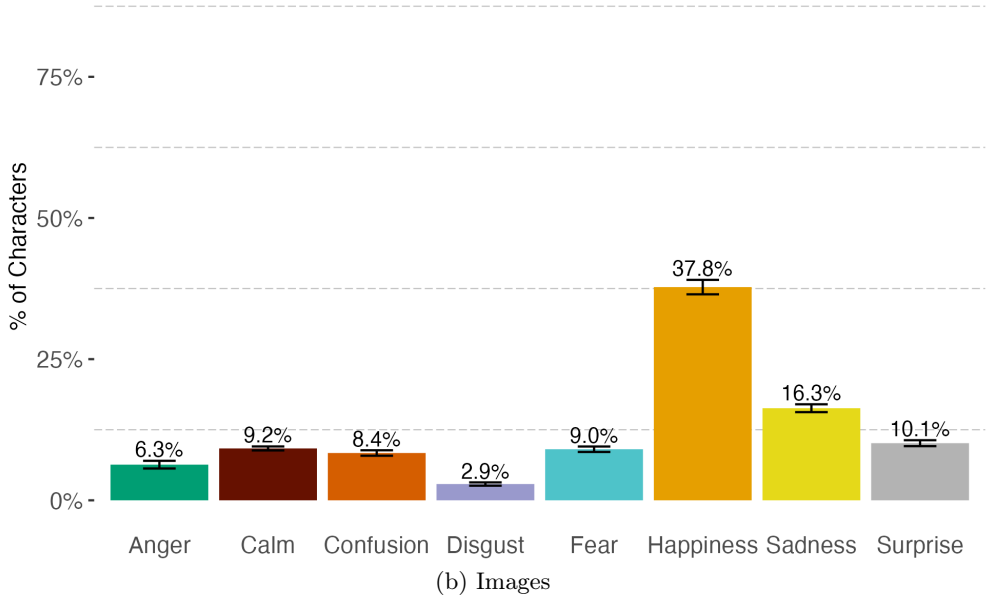
(d) Images: Children's Books



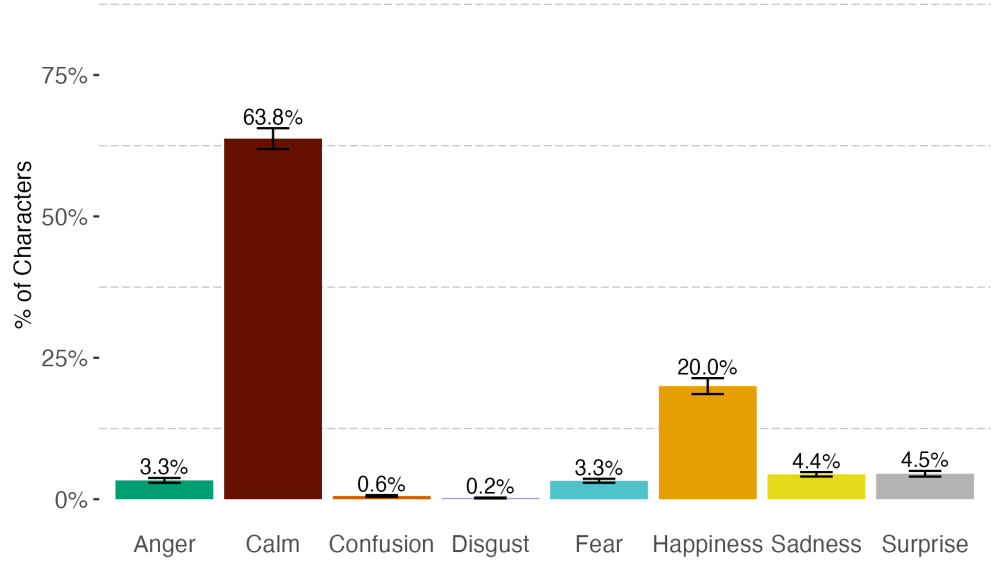
Notes: This figure replicates Figure 3 with additional categories for "No Emotion" and "Unsure." Panel (a) shows the share of all emotions detected in the sample of textbooks, and panel (b) shows the shares in literature. Emotions are detected using OpenAI. Emotions for each textbook are counted for each year the textbook is in use, and emotions for each children's book are counted for the book's copyright year. Panels (c) and (d) show the analogous results for emotions detected in faces in both corpora. Children's books are associated with their respective copyright years, whereas textbooks are associated with each year they are in use. Note that textbooks are in use for multiple years, whereas the children's books only appear one time in the time series panel.

FIGURE A.3
Overall Representation of Emotions (Matched Sample)

(a) Text

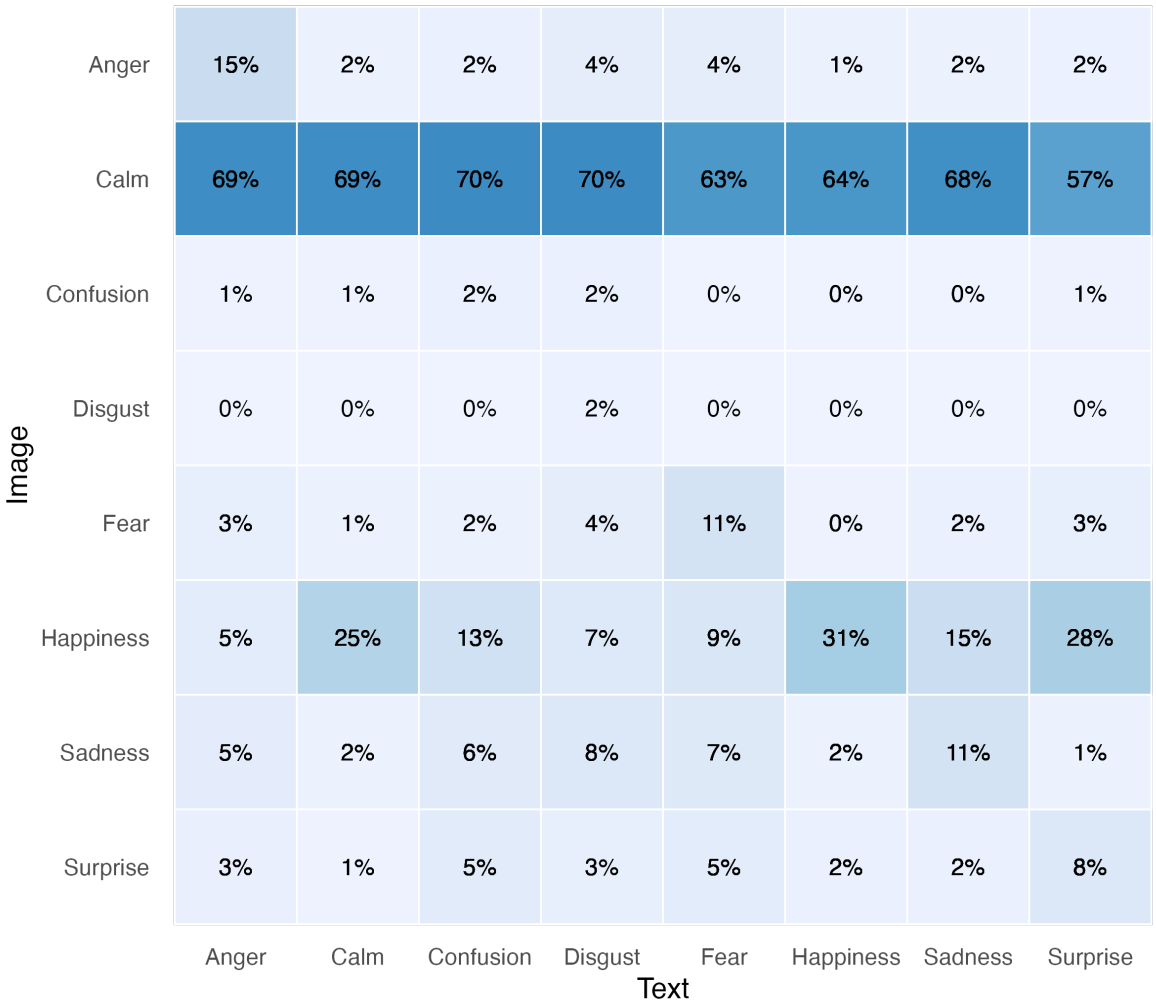


(b) Images



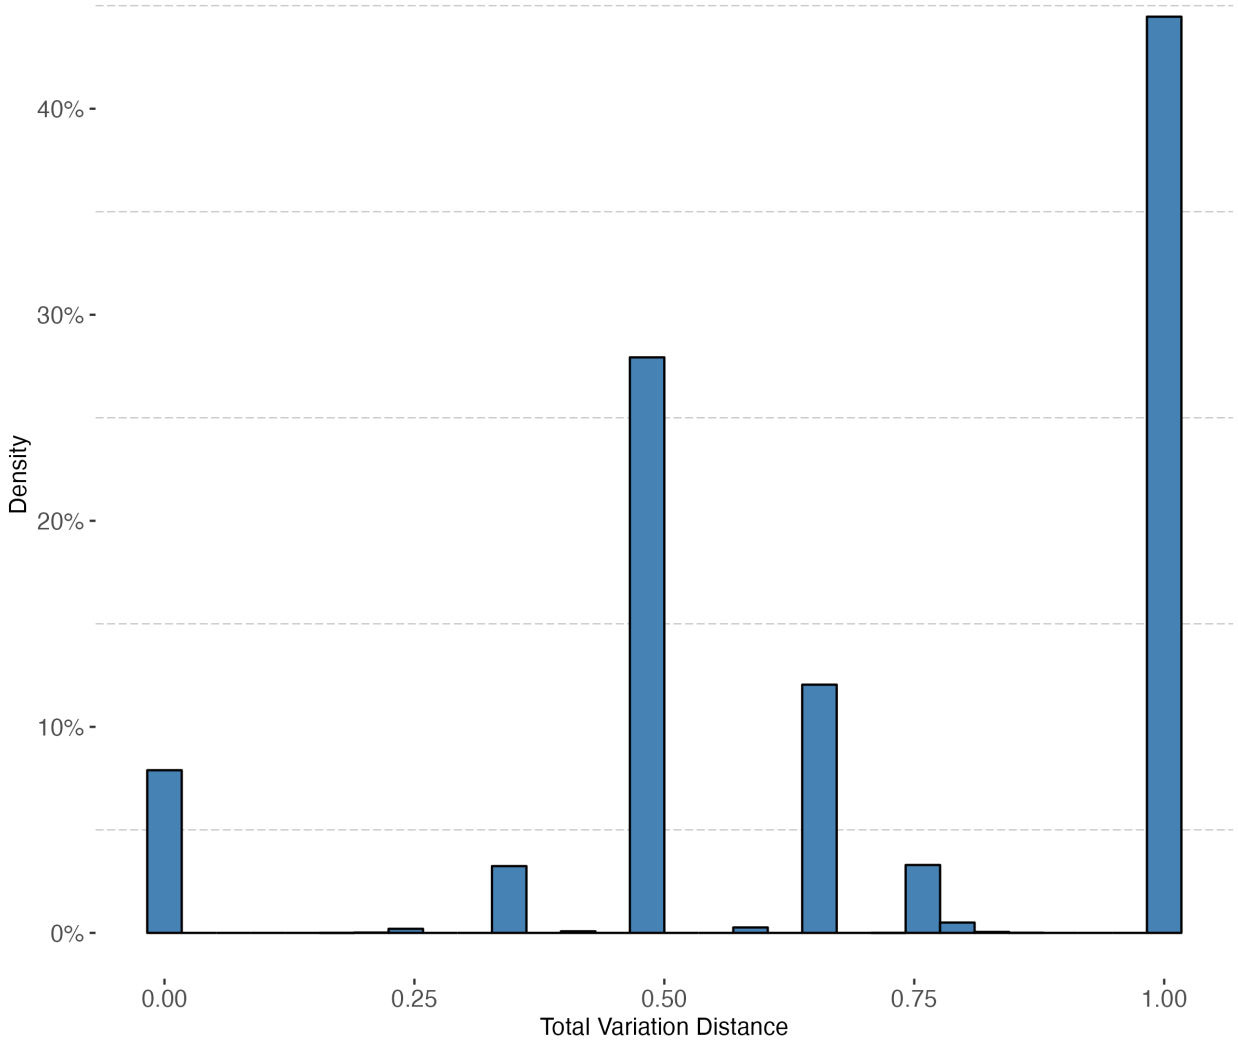
Notes: This figure shows the aggregate displays of emotions in text and images for the sample of matched pages. The top panel shows the percentage of emotions conveyed by characters as a share of all emotions detected in the sample of books (literature and textbooks); emotions are detected using an LMM. The bottom panel shows the percentage of emotions conveyed by faces as a share of all emotions detected in the sample of books (literature and textbooks); emotions are detected using an LMM.

FIGURE A.4
Normalized Co-occurrence Matrix for Characters Matched in Text and Images



Note: This figure shows the co-occurrence of emotions for text and images on the same page, where unique human characters are matched on their race and gender. Co-occurrence is normalized by the text columns. In other words, each cell shows the proportion of times a specific image label occurred, given a particular text label. For example, the top left cell shows that for pages where there is one Black male character in the text and the image, when the character is shown as angry in text,the matched character will be depicted as angry in an image only 15% of the time.

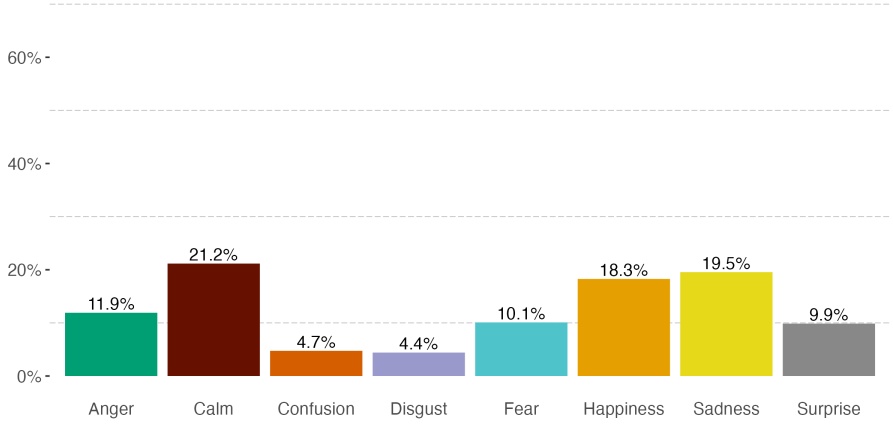
FIGURE A.5
Total Variation Distance Between Text and Images



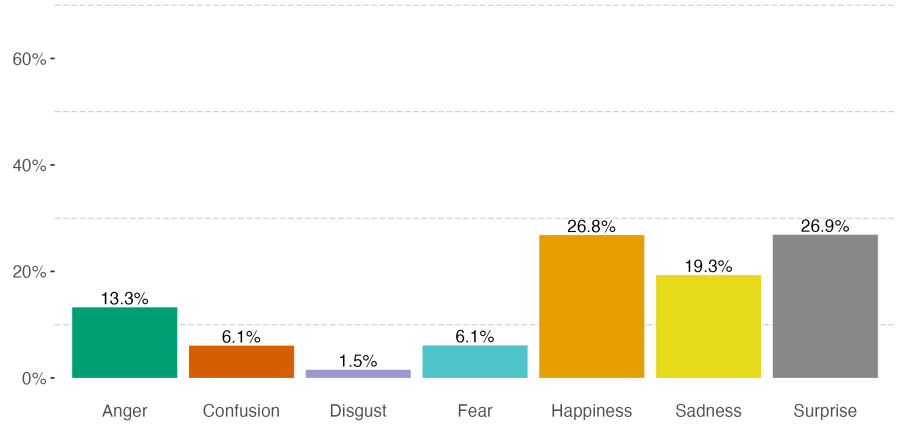
Note: This figures shows the total variation distance (tvd) between the emotional content of images and text on the same page. We calculate tvd by taking the l^1 distance between the empirical distributions of our 8 emotions in text and in images on a given page and multiply by $1/2$.

FIGURE A.6
Robustness: Overall Representation of Emotions

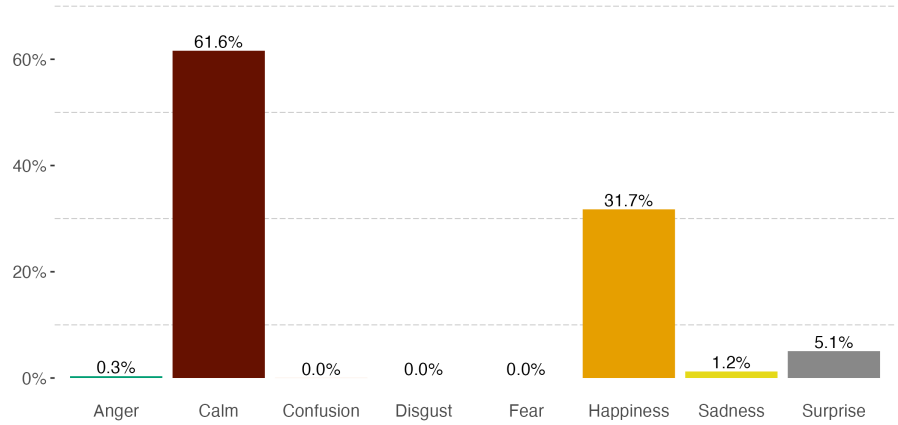
(a) Text: Lexicon Approach



(b) Text: GoEmotions AI Model

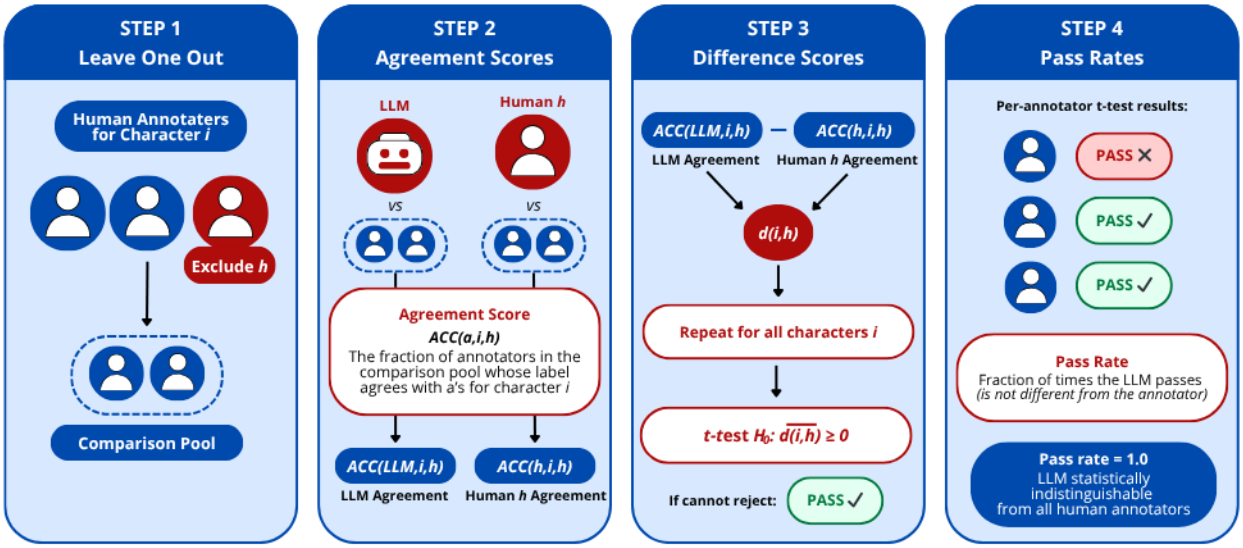


(c) Images: Feature Classification



Notes: This figure shows the aggregate displays of emotions in text and images using alternative approaches. Panel (a) shows the percentage of emotion-specific words as a share of all emotion words from a pre-specified lexicon. Panel (b) shows the percentage of emotion-specific sentences as a share of all emotion sentences as classified by the GoEmotions model. Panel (c) shows the percentage of detected faces.

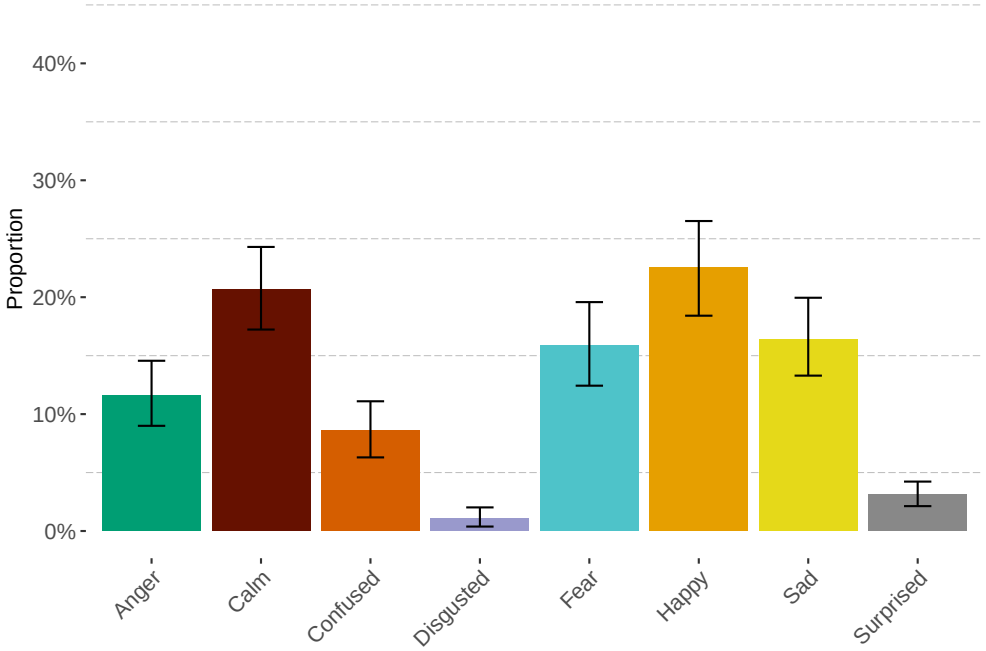
FIGURE A.7
Alternative Annotator Test Procedure



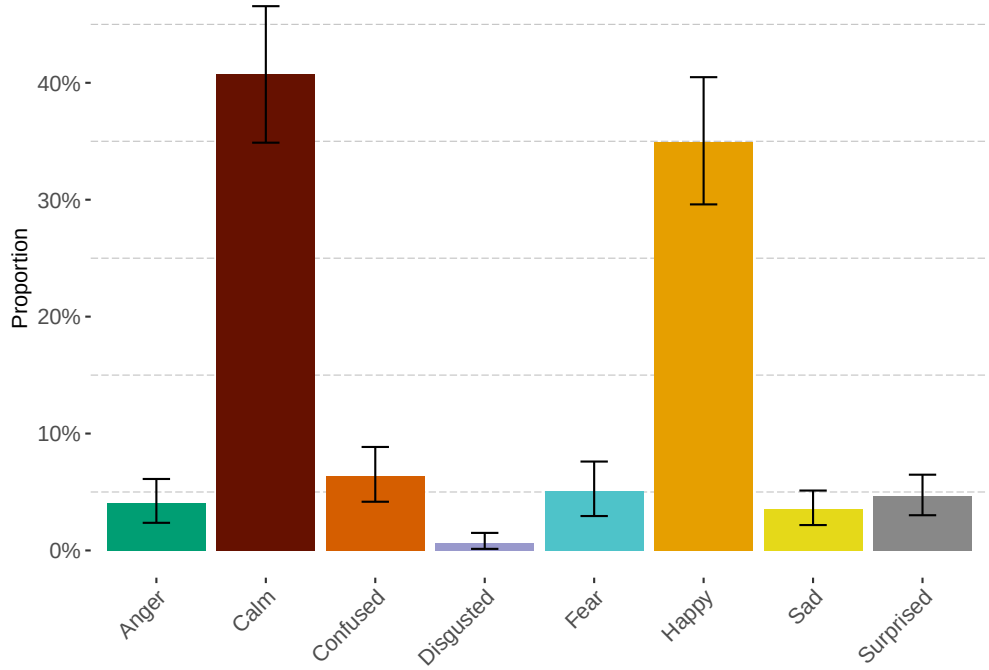
Note: This figures shows the procedure used to validate the AI output. The process involves assessing the model's reliability relative to human annotators in terms of its relative capacity to agree with other raters on the same set of examples.

FIGURE A.8
Simulated Distribution of Human Labels

(a) Text



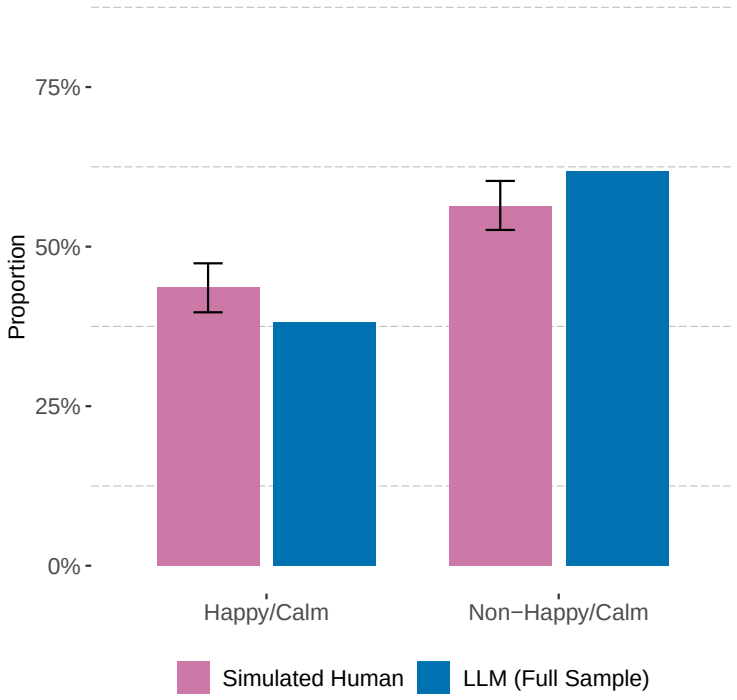
(b) Images



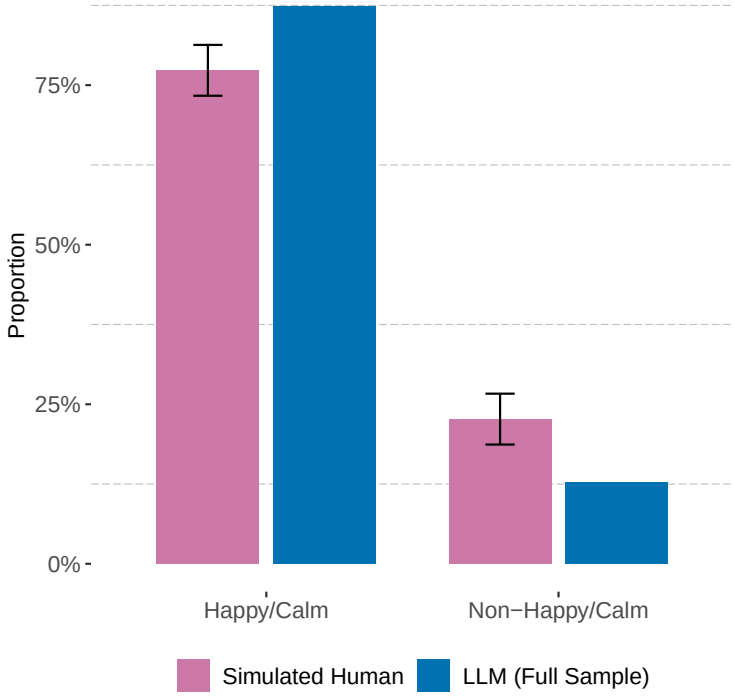
Notes: This figure shows the distribution of emotions predicted to be obtained by human labelers on the full sample. The simulation procedure is described in Section V.B. The top panel shows the distribution in text, and the bottom shows the distribution in images.

FIGURE A.9
Simulated Human Labels versus Models on Happy/Calm and Non-Happy/Calm

(a) Text

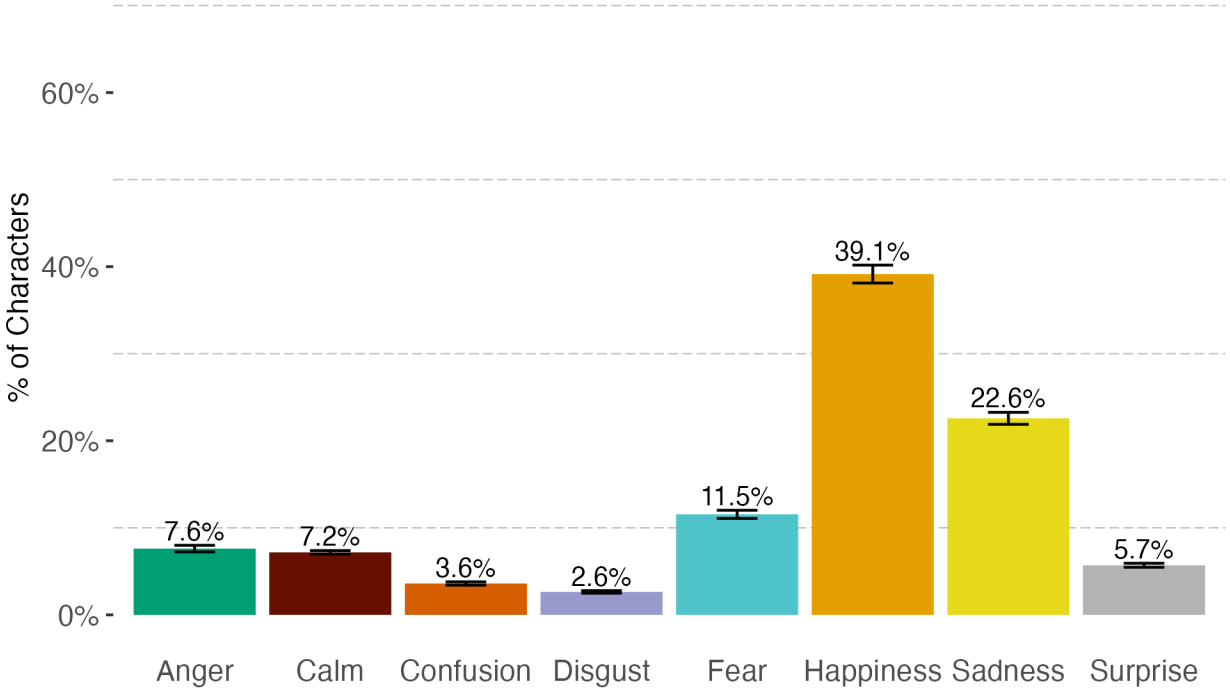


(b) Images



Notes: This figure compares the human and model predictions on the Happy/Calm and Non-Happy/Calm split. The top panel shows the distribution in text, and the bottom shows the distribution in images.

FIGURE A.10
Overall Representation of High Confidence Emotions In Text



Note: This figure shows the aggregate displays of emotions in text when we subset to emotion labels that the model has high confidence in (≥ 0.8). Error bars depict 95% confidence intervals constructed from standard errors that are clustered at the book level.

B Appendix Tables

TABLE B.1
Sample of Mapping from Keyword Lexicon to Emotions in Main Analysis

anger	**calm**	**confusion**	**disgust**	**fear**	**happiness**	**sadness**	**surprise**
acerb	abide	absentminde	abhor	affright	affection	abandon	aah
acerbic	amiable	absurd	abhorrence	afraid	amuse	abandonment	aflutter
acerbity	aplomb	absurdity	abhorrent	aghast	amusing	ache	ah
acrid	asleep	addle	abhorringly	alarm	banter	afflict	aha
acrimonious	assuage	baffle	abominable	alarum	beam	affliction	ahha
...	...	...	...	...	...	...	...

Note: This table presents a sample of our lexicon dictionary. Each of the eight emotion categories is associated with a mutually exclusive set of keywords. The full lexicon can be viewed here

TABLE B.2
Mapping from GoEmotions to Emotions in Main Analysis

anger	**confusion**	**disgust**	**fear**	**happiness**	**sadness**	**surprise**
anger	confusion	disgust	fear	amusement	embarrassment	curiosity
annoyance			nervousness	excitement	disappointment	realization
disapproval				joy	grief	surprise
				love	remorse	
				optimism	sadness	

Note: This table presents the mapping used in our main analysis to group the 28 fine-grained GoEmotions into the eight broader emotion categories used in our main analysis. We exclude the following fine-grained emotions from the joy category: admiration, approval, caring, desire, gratitude, pride, relief.

TABLE B.3
Macro F1 Scores

Prediction	**Macro F1**	**Number of Examples**	**Exclusion Rate**
Panel A: Text			
Primary Emotion	0.289	588	35.0%
Emotion (Binarized)	0.467	588	11.7%
Is Human	0.862	564	0.7%
Gender	0.802	554	3.4%
Age	0.680	588	7.8%
Race	0.531	532	7.3%
Panel B: Image			
Primary Emotion	0.480	459	22.9%
Emotion (Binarized)	0.640	459	5.7%
Is Human	0.941	459	0.9%
Gender	0.821	459	2.6%
Age	0.782	368	3.8%
Race	0.685	305	12.5%
Skin Color	0.636	333	10.2%

Notes: This table reports macro F1 scores, which are obtained by computing the F1 score for each class and averaging over classes. The number of examples changes because raters may find demographic characteristics to be not applicable (i.e., the race of a non-human character). Examples are excluded if all three human raters disagree so that there is no modal rater. The high exclusion rate for text and image emotion labeling reflects the relative difficulty of this task.

TABLE B.4
Alternative Annotator Test

Prediction	**Avg Advantage**	**Pass Rate** ($q = 0.10$)	**Pass Rate** ($q = 0.05$)
Panel A: Text			
Primary Emotion	−9.3%	0.333	0.333
Emotion (Binarized)	−8.5%	0.333	0.333
Is Human	−0.8%	1.000	1.000
Gender	5.1%	1.000	1.000
Age	5.2%	1.000	1.000
Race	−11.3%	0.333	0.333
Panel B: Image			
Primary Emotion	−3.9%	0.500	0.833
Emotion (Binarized)	−6.3%	0.167	0.667
Is Human	2.3%	1.000	1.000
Gender	−1.1%	1.000	1.000
Age	−7.2%	0.333	0.500
Race	1.3%	1.000	1.000
Skin Color	−4.8%	0.500	0.500

Notes: This table reports the results from the alternative annotator test. The top panel shows the results for the text predictions, and the bottom panel shows the results for the image predictions. Binarized emotion is constructed from the primary emotion label by binning into Happy/Calm and Non-Happy/Calm categories (along with a No Emotion/Unsure category).

C Methods Appendix

C.A Text Analysis

LMM Prompt Used for Emotion Identification Our primary text model is OpenAI's GPT-4o-mini, version 2024-07-18, accessed through the API in April 2025 with temperature set to 1.0 and the system prompt: "You are a knowledgeable and unbiased judge." For each page of OCR-extracted text, we submitted a standardized prompt which instructed the model to classify each character along with their age, gender, race, whether or not they were human, and any emotions expressed by that character. Emotion labels were recorded under three taxonomies: the six Ekman emotions, our eight-category MiiE taxonomy, and the 28-category GoEmotions taxonomy, each with a confidence score. The prompt restricted responses to predefined answer sets, allowed "unsure" and "not applicable" values for ambiguous cases, required "no emotion" for characters without expressed affect, and instructed the model not to assign emotions or attributes unsupported by the text. Hallucinations of emotions beyond simple morphological variants account for less than 1% of all observations, indicating high prompt adherence. The full prompt can be viewed at https://github.com/miielab/emotions_paper. This gives us a final dataset containing character-page level emotions.

Alternative Text Analysis Approach: Using a Lexicon-Based Approach To create our lexicon, we primarily use the synsets of WordNet. A synset is a collection of cognitive synonyms that convey the same concept. Some words have a single synset, while others have multiple. For example, the word "anger" is part of the following synsets: 1) a feeling that is oriented toward some real or supposed grievance; 2) the state of being angry; and 3) belligerence aroused by a real or supposed wrong (personified as one of the deadly sins). Hence, with an initial emotion-related search term, we can access a comprehensive network of semantic relations.

We constructed an initial set of WordNet inputs, extract their synonyms, and manually review the results. Following this, we engineer a prompt to elicit more emotional synonyms from an LLM. Our final lexicon maps these synonyms to our 8 categorical emotions (confusion, anger, sadness, surprise, calm, happiness, fear, and disgust). In the process, we lemmatize all words using SpaCy and, through manual inspection, ensure that all categories are mutually exclusive and each keyword unique. Lemmatization is chosen in order to simplify the lexicon construction at the cost of less flexibility over word choice. Table B.1 shows a sample of the dictionary. The full lexicon can be viewed at https://github.com/miielab/emotions_paper.

In total, our lexicon contains 1,706 entries. More specifically, breaking this down into the number of synonyms for each emotion, we have 253 anger, 256 calm, 191 confusion, 224 disgust, 173 fear, 204 happiness, 246 sadness, and 159 surprise entries.

The sample of textbooks and children's literature is then processed by splitting it into sentences using SpaCy and then cleaned by removing certain punctuation and digits, by making it lowercase and by lemmatizing it. This gives us 3,361,503 sentences across 1,306 books. Subsequently, we apply a count vectorizer to each sentence, allowing us to identify if a sentence contains one or more keywords and the number of times it appears. This gives us a final dataset containing sentence level emotions.

Alternative Approach: Using an Emotions AI-Based Approach This approach uses a RoBERTa-based model trained on the GoEmotions dataset created by Demszky et al. (2020). It is a multi-label classification model that takes a document as an input and returns 28 'probability' float outputs, one for each of the 28 emotion labels.

We use the same sample of 3,361,503 cleaned sentences as in the lexicon approach, but in this case do not lemmatize. We then feed these sentences into the model, receive the output, and assign each sentence the emotion with the highest predicted probability. Sentences labeled as "neutral" are dropped, as we are only concerned with text that has emotional content. This parallels our lexicon method, where our denominator consists only of emotion-related words. Next, we map each of the 27 emotions to 7 aggregate emotions guided by a correspondence released by Demszky et al. (2020). Ideally, we would like to have aligned the output to our MiiE Lab 8, but the model outputs did not allow us to easily identify "calm." This gives us a final dataset containing sentence level emotions.

In comparison to the lexicon-based approach, this method takes into account the entire context of sentence when making a prediction. The base model was pretrained on a large volume of texts and learned the contextual representation of words based on their surroundings. Hence, it is able to reduce semantic ambiguity otherwise introduced by analyzing each word individually. However, we lose some of the transparency our lexicon-based method provides through insight into the words with which emotions are being conveyed.

C.B Image Analysis

We quantify emotional representation in illustrated pages of children's books using a two-stage, fully automated pipeline. We implement the pipeline on Google Cloud Vertex AI using Gemini 2.5 Pro Preview, model version `gemini-2.5-pro-preview-03-25`.

First, each scanned page is processed using a lightweight Gemini classifier. This classifier assigns a binary label (either `TEXT_ONLY` or `CONTAINS_IMAGE`) using a low temperature setting of 0.1 and a response limit of 256 tokens to ensure consistent binary classification. Pages identified as `TEXT_ONLY` are excluded from further analysis, whereas pages labeled `CONTAINS_IMAGE` pages advance to the next stage in the pipeline.

Next, for our illustrated pages, the model outputs character level data (e.g., character bounding boxes, demographic attributes, and emotion labels), using a temperature 0.0, a 10,000-token limit. Importantly, the character level output includes three distinct emotion codings (`emotion_6_options`, `emotion_10_options`, `emotion_28_options`).

Because the image prompt allowed for open-ended race predictions, we standardized the model output before using it in the analysis. In particular, the re-categorized variable for race (race_original) was constructed according to the following classification of the Gemini model's output.

- Not applicable: Occurs when the character is classified as non-human
- Black: if race_original contains any of: black, african, afro, ashanti, mende
- Latine: if race_original contains any of: latin, cuban, rican, hispanic, mexica, mestizo, caboclo, mestiza
- Asian: if race_original contains any of: asian, afghan, chinese, hmong, indian, japanese, korean, mongol, tibetan, vietnamese, cambodian, filipino, kashmiri, kurdish, pashtun
- White: if race_original contains any of: white, greek, egyptian, european, german, russia, spanish, armenian, balkan, hungarian, italian, kosovan, mediterranean, norse, roman, celtic, jewish, ashkenazi
- Indigenous if race_original contains any of: indigenous, native, inuit, aboriginal, aben, aztec, guarani, hawaiian, inuk, inupiaq, kwakiutl, maya, navajo, olmec, penan, polynesian, saami,

samoan, sioux, yup, zapotec, pacific islander, melanesian

- Unsure: if race_original contains any of: unclear, unknown, unspecified, unsure, ambiguous
- Other: if race_original contains any other classification