



Educator Attention: How computational tools can systematically identify the distribution of a key resource for students

Qingyang Zhang
Stanford University

Rose E. Wang
Stanford University

Ana T. Ribeiro
Stanford University

Dorottya Demszky
Stanford University

Susanna Loeb
Stanford University

Educator attention is critical for student success, yet how educators distribute their attention across students remains poorly understood due to data and methodological constraints. This study presents the **first largescale computational analysis of educator attention patterns**, leveraging over 1 million educator utterances from virtual group tutoring sessions linked to detailed student demographic and academic achievement data. Using natural language processing techniques, we systematically examine the recipient and nature of educator attention. Our findings reveal that educators often provide more attention to lower-achieving students. However, disparities emerge across demographic lines, particularly by gender. Girls tend to receive less attention when paired with boys, even when they are the lower achieving student in the group. Lowerachieving female students in mixed-gender pairs receive significantly less attention than their higher-achieving male peers, while lower-achieving male students receive significantly and substantially more attention than their higher-achieving female peers. We also find some differences by race and English learner (EL) status, with low-achieving Black students receiving additional attention only when paired with another Black student but not when paired with a non-Black peer. In contrast, higher-achieving EL students receive disproportionately more attention than their lower-achieving EL peers. This work highlights how large-scale interaction data and computational methods can uncover subtle but meaningful disparities in teaching practices, providing empirical insights to inform more equitable and effective educational strategies.

VERSION: March 2025

Suggested citation: Zhang, Qingyang, Rose E. Wang, Ana T. Ribeiro, Dorottya Demszky, and Susanna Loeb. (2025). Educator Attention: How computational tools can systematically identify the distribution of a key resource for students. (EdWorkingPaper: 25 -1144). Retrieved from Annenberg Institute at Brown University: <https://doi.org/10.26300/hvq6-d385>

Educator Attention: How computational tools can systematically identify the distribution of a key resource for students

Qingyang Zhang*, Rose E. Wang*
Ana T. Ribeiro, Dora Demszky^α, Susanna Loeb^α

Stanford University
q1zhang@stanford.edu, rewang@cs.stanford.edu

Abstract

Educator attention is critical for student success, yet how educators distribute their attention across students remains poorly understood due to data and methodological constraints. This study presents the **first large-scale computational analysis of educator attention patterns**, leveraging over 1 million educator utterances from virtual group tutoring sessions linked to detailed student demographic and academic achievement data. Using natural language processing techniques, we systematically examine the recipient and nature of educator attention. Our findings reveal that educators often provide more attention to lower-achieving students. However, disparities emerge across demographic lines, particularly by gender. Girls tend to receive less attention when paired with boys, even when they are the lower achieving student in the group. Lower-achieving female students in mixed-gender pairs receive significantly less attention than their higher-achieving male peers, while lower-achieving male students receive significantly and substantially more attention than their higher-achieving female peers. We also find some differences by race and English learner (EL) status, with low-achieving Black students receiving additional attention only when paired with another Black student but not when paired with a non-Black peer. In contrast, higher-achieving EL students receive disproportionately more attention than their lower-achieving EL peers. This work highlights how large-scale interaction data and computational methods can uncover subtle but meaningful disparities in teaching practices, providing empirical insights to inform more equitable and effective educational strategies.

1 Introduction

Educator attention is one of the most valuable resources for building effective learning experiences—whether it’s in the classroom, extra-curricular activities, or tutoring. When students receive more attention from educators, such as instructional or emotional support, students are more likely to learn and acquire the skills they need to engage successfully in school (Birch and Ladd, 1997, Burchinal et al., 2008, Cortes et al., 2024, Hamre and Pianta, 2005, Nickow et al., 2024). Recognizing the important role of educator’s attention to student-specific needs, post-pandemic education policies have prioritized increasing the amount of attention students receive, such as by expanding tutoring programs (FutureEd, 2022, Groom-Thomas et al., 2023).

However, the presence of an educator does not guarantee that all students receive the same kind of attention. Prior work has documented differences in how educators allocate attention, with student demographics influencing the quantity and nature of the attention students receive. Female students and students from racial minority groups, for instance, often receive less attention or lower quality instruction, such as being given easier content (Beaman et al., 2006, McKown and Weinstein, 2008,

* = Equal contributions. ^α = Equal advising.

Watson et al., 2022). Educator attention is therefore a double-edged sword: While it has the potential to remediate education inequalities, it can also inadvertently perpetuate them.

Understanding the patterns in educator attention requires large-scale, systematic analysis of both whom educators address and how they engage with students. Traditionally, researchers have relied on classroom observations where trained annotators manually track educator-student interactions, such as how often an educator calls on students (Allen et al., 2013, Brown et al., 2010, Gunn et al., 2021, Younger et al., 1999). While valuable, these observation methods are limited in scale, capturing only some hours of instruction, and lack systematic control over a student’s academic background that shape educator attention beyond a student’s demographics (Reinholz and Shah, 2018). For example, if educators interact more frequently with male students, is this due to gender alone or due to male students struggling more academically than female students?

The recent expansion of virtual tutoring, where educators meet their students over video, has created unprecedented opportunities to study educator-student interactions at scale (Cortes et al., 2025, Gortazar et al., 2024). Each interaction between educators and their students is recorded and linked with individual student characteristics. Coupled with advances in natural language processing (NLP), this data allows researchers to systematically analyze both the amount and nature of educator attention. Researchers have used NLP across domains to examine large-scale language patterns in high-stakes interactions, including police interactions (Voigt et al., 2017), classroom interactions (Jacobs et al., 2022, Suresh et al., 2022) and tutoring interactions (Demszky et al., 2024, Wang et al., 2024a).

In this work, we analyze over 1,157,970 educator utterances in tutoring sessions and investigate how educator attention varies by student achievement and demographics. We focus on three key demographic indicators studied in prior work on educator attention: gender, race and English Learner (EL) status. Our data comes from a randomized controlled trial of a tutoring program found to be effective at improving student literacy skills (Robinson et al., 2024). Using NLP techniques, we automatically identify who educators direct their attention to and how they engage with the students.

Our analysis provides the first large-scale quantitative study of educator attention patterns, where we both validate and extend previous qualitative research. Concretely, we find that educators provide more attention to lower-achieving students, suggesting a responsiveness to student needs. However, we also reveal notable deviations from this pattern: Lower-achieving female students in mixed-gender pairs receive significantly less attention than their higher-achieving male peers, whereas lower-achieving male students receive substantially more attention than their higher-achieving female peers. A similar pattern arises for Black students: Low-achieving Black students receive more attention than their peer only when paired with another Black student but not when paired with a non-Black peer. By contrast, higher-achieving EL students receive significantly more attention than their lower-achieving EL peers. These findings suggest that even within effective tutoring programs, student demographics and group compositions may unconsciously shape educator attention. Our study demonstrates how large-scale analysis of fine-grained educator-student interactions can uncover subtle but meaningful disparities, and surface opportunities to promote better teaching practices.

Data

Our dataset consists of transcribed video recordings from a U.S.-based virtual tutoring program focused on early literacy for K-2 students, which was found to improve student achievement scores

		NATURE OF ATTENTION		
		Content	Relationship-Building	Management
RECIPIENT OF ATTENTION	Only Student A	"Okay, [Student A], can you use this word in a sentence, please?"	"Okay, [Student A], tell me one fun thing from the weekend."	"Oh, [Student A], I can't see you."
	Only Student B	"Okay, [Student B], what's the middle sound?"	"I know [Student B] likes to read."	"And [Student B], can you hit mute?"
	Both Students	"Ooh, mushroom does have the M sound."	"You guys are awesome."	"Let's keep our listening ears on and our focus, and let's do this page."
	One of the Students	"Ball."	"Today you're six years old?"	"Okay, keep your earphones on, stop moving the computer screen."

Table 1: **A sample of utterances illustrating our two-dimensional Attention Framework.** Our framework classifies a teacher utterance based on both the NATURE OF ATTENTION (content, relationship-building, or management) and the RECIPIENT OF ATTENTION. The classification of each utterance relies on the broader conversational context (not shown in the table). While some utterances can be clearly categorized, many require careful consideration of context for accurate classification.

in a randomized controlled trial (Robinson et al., 2024). This dataset includes 1,157,970 utterances from 5,249 2-on-1 tutoring sessions, where a tutor teaches two students in a session. Students in each pair were matched at the beginning of the school year, and largely remained in the same pair throughout the year for consistency. The tutors are based in the U.S., and are mostly former classroom teachers and part-time teachers. Our dataset includes 125 educators and 793 students.

Every tutoring session is linked to student information provided by the school district, including gender, race, EL status, and achievement. Given prior work on educator attention, we focus on three key demographic factors frequently studied in this literature, gender (Beaman et al., 2006), race (Reinholz and Shah, 2018), and EL status (Solano-Flores et al., 2024). We also account for students' pre-intervention achievement measured by their Dynamic Indicators of Basic Literacy Skills (DIBELS) scores, a common measurement of student literacy skills (Good III et al., 2001). Please refer to Appendix Section A for more details on data.

Measuring Attention

Traditional classroom observations are effective in discovering patterns in educator attention, however are too resource-intensive for large-scale analysis. Scalable, systematic approaches are important for advancing research on how educator attention varies in different settings and where to improve teaching practices.

In this study, we develop computational classification models of the recipient and of the nature of attention, and tune them on 1,200 labeled utterances with preceding context from the interaction. Our classification taxonomy, shown in Table 1, synthesizes theoretical frameworks from prior education research on instructional attention with empirical observations from our dataset. This approach ensures our measures are both theoretically grounded and practically relevant. See Materials and Methods for more detail.

By applying these trained models to our complete dataset of 1,157,970 utterances, we conduct a systematic inquiry of educator attention patterns and their relationship to student characteristics and academic needs. Our work focuses on the educator’s utterances, and we do not model student-initiated behaviors that can also shape how educators allocate attention (Brophy and Good, 1970, Irvine, 1986). Our findings characterize the distribution of educator attention and do not establish causal relationships.

Recipient of Attention

Previous studies have documented differences in educator attention across students. For example, research has shown that classroom teachers give less individual attention to female students compared to male students (Beaman et al., 2006, Heller and Parsons, 1981, Jones and Dindia, 2004). Studies of race differences have produced mixed results and research focusing on differential attention by EL status is sparse (Cooper and Allen, 1998, Irvine, 1986, Solano-Flores et al., 2024). Our computational approach provides the opportunity for us to explore differences across gender, race and EL status in educator attention by identifying the recipient of that attention leveraging past work on addressee identification (Jovanovic and op den Akker, 2004, Ouchi and Tsuboi, 2016). In our two-student tutoring context, we classify each educator utterance into four categories: directed to student A, directed to student B, directed to both students, or directed to an ambiguous recipient. With this classification, we can quantify the amount of attention given to a student by measuring the total duration of utterances directed to that student. Additionally, by computationally measuring the amount of attention to a student, we can analyze how educator attention varies across student characteristics at a scale previously impossible with manual methods. Our main study focuses on the attention directed to individual students, and Appendix Section B reports results on attention directed to both students.

Nature of Attention

Beyond measuring who receives attention, prior work has studied the differences in the nature of attention. For example, researchers have found that boys receive more attention about their behaviors than girls do (Dweck et al., 1978, Irvine, 1986). Prior work has also documented the importance of relationship-building and content-focused attention, because these forms of attention can help build student confidence and improve achievement especially for historically under-performing student groups including EL students (Lewis et al., 2012, Soland and Sandilos, 2021, Soland, 2019). Mirroring established classroom observation rubrics (Pianta, 2008) and drawing on empirically grounded observations, we develop a classification taxonomy for teacher attention that captures three key dimensions of instructional practice: Content-focused instruction (discussing academic material), relationship-building (developing rapport and engagement), and session management (coordinating student behaviors and activities).

Findings

Study 1: Achievement-Based Patterns

Prior research provides evidence that educator attention can effectively bolster learning for lower-achieving students (Lewis et al., 2012, Soland and Sandilos, 2021). As a result, educators may consider students’ achievement levels when determining how to allocate their attention across students and particularly focus on students who struggle with the material (Brophy and Good, 1970, Myhill, 2002). We test whether educators provide more attention to lower-achieving students. In our paired tutoring context, we define the lower-achieving student as the one with the lower baseline achievement score; see Appendix Section A on descriptive statistics of lower- and higher-achieving

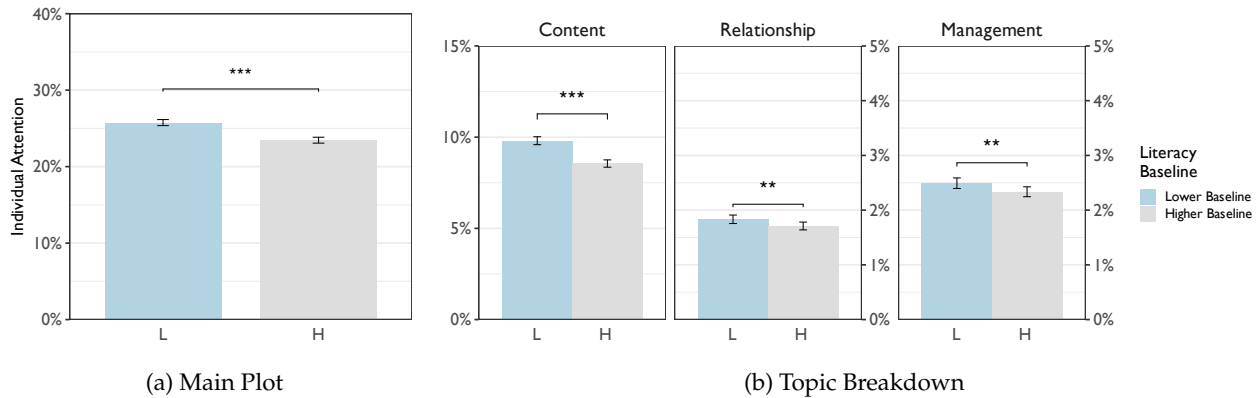


Figure 1: **Individual attention allocation based on student achievement.** (a) Lower-achieving students (L) receive significantly more individual attention than their higher-achieving peers (H). (b) Breakdown of attention by type, showing that lower-achieving students receive more attention on content-related support, relationship-building interactions, and session management. Asterisks indicate statistical significance (** $p < 0.01$, *** $p < 0.001$).

students. Figure 1 reports the distribution of attention given to the lower-achieving and higher-achieving student within a tutoring pair, in terms of both (a) overall attention, and (b) the nature of attention.

We find that educators are more attentive to students who have greater learning needs, with additional 2.3 percentage points (p.p.) attention to the lower-achieving student ($t(5248) = 7.63$, $p < 0.001$). This pattern holds across the types of attention with 1.25 p.p. for content, 0.12 p.p. for relationship and 0.16 p.p. for management (Figure 1b). The most pronounced category is content-focused attention given to the lower-achieving student ($t(5248) = 8.99$, $p < 0.001$). Overall, Study 1 suggests that educators are responsive to students' academic needs, providing them especially with more attention on instructional content.

Study 2: Demographic-Based Patterns

In addition to academic achievement, students' demographics—such as gender, race and home language—may also affect the attention they receive from educators. For example, Gilliam et al. (2016) conducted an eye-tracking study on preschool teachers and found that they tended to focus more on Black boys when asked to look for potential behavioral issues, even when no misbehavior occurred. Similarly, the demographic composition of student pairs—such as whether students are in same-gender or mixed-gender pairs—may further influence how educators distribute attention (Kulik and Kulik, 1982). While some differences in attention may stem from educators prioritizing lower-achieving students, others may persist beyond achievement and potentially reinforce structural disparities in educational opportunities.

To examine this in Study 2, we analyze how tutors allocate attention based on individual student demographics and the demographic composition of student pairs, while controlling for academic achievement. We focus on three key demographic indicators studied in prior work on educator attention: gender, race and EL status, as defined by the district and available in district administrative data. For gender, we use available classifications of male or female. For EL status, we compare EL and non-EL students. For race/ethnicity, we categorize students as Black or non-Black; while we have a more detailed measure of race and ethnicity, the sample size for non-Black, non-Hispanic




Figure 2: **Regression coefficients for individual attention allocation across student pairings.** The left-most dot is the reference group. Error bars indicate standard errors. (a) Gender pairings: In mixed-gender groups, female students receive significantly less attention than their male peers (the reference group). (b) Race pairings: In groups with different races with a focus on Black students, Black students receive slightly more attention than their non-Black peer (the reference group). (c) EL status pairings: In groups with students of different EL status, we find that students get about the same amount of attention.

students is too small to reliably estimate differences between each subgroup (see Appendix Section A for a sample breakdown). We focus on Black students because of the especially strong systemic barriers they often face, including economic disadvantage, lower test scores and low graduation rates (Chetty et al., 2020, for Educational Statistics et al., 2024).

Concretely, we model the amount of attention, y_{ijk} , a student i receives when partnered with student j in session k as a function of: an indicator of their pair composition I_{ijk} (e.g., female in same-gender or mixed-gender pairings), their own baseline achievement D_i , their partner’s achievement D_j and their relative achievement (lower or higher) R_{ij} . We cluster standard errors at the pair level, e_{ij} .

$$y_{ijk} = \beta_1 I_{ijk} + \beta_2 D_i + \beta_3 D_j + \beta_4 R_{ij} + e_{ij}$$

Figure 2 reports β_1 , the estimated differences in educator attention based on the student-level (gender, race, and EL status) and group demographics (same or mixed gender, race, and EL). We also report differences by the nature of attention below and include the plots in Appendix Section B.

Figure 2a reports the amount of attention students of different genders and gender pairings receive. Our findings highlight a pronounced gender gap in attention allocation. Female students receive significantly less individual attention than their male counterparts, with the largest disparity emerging in mixed-gender pairings, where female students receive 5.2 p.p. less attention than their male partners [$\beta = -0.052(-0.081, -0.023)$, $p < 0.001$], even when controlling for achievement levels. Female students particularly receive less attention on content and management (rf. Appendix Section B).

Figure 2b reports the results for race, and Figure 2c the results for EL status. We find that when both students are non-Black or both are non-EL students, they receive significantly less individual attention compared to pairs where at least one student is Black or is an English learner. Specifically,

pairs of non-Black students receive 5.9 p.p. less attention [$\beta = -0.059(-0.082, -0.037)$, $p < 0.001$], and pairs of non-EL students receive 5.7 p.p. less attention [$\beta = -0.057(-0.081, -0.034)$, $p < 0.001$] compared to non-Black students and non-EL students in mixed pairings, respectively. This reduction is primarily driven by a overall decrease in content-focused attention (rf. Appendix Section B). Study 2 thus suggests that educators allocate disparate attention to students based on their demographics, beyond what is explained by students' measured achievement.

Study 3: Demographic Patterns by Achievement Relative to Peers

Studies 1 and 2 established two key findings: Educators generally prioritize the lower-achieving student in the pair (Study 1), but they also show systematic differences in attention allocation based on student demographics (Study 2). These findings raise a critical question: Does the tendency of educators to support the lower-achieving student of the group operate uniformly across demographics?

To address this question, we examine whether a student's relative position as the lower- or higher-achieving student of their pair influences the attention they receive based on their demographics. For instance, when a Black student is the lower-achieving student of the pair, does the student receive the additional attention we would expect based on Study 1's findings? We report our results in Figure 3 where we compare the amount of individual attention given to students when they are the lower- versus higher-achieving student of their pair, across different demographic groups. We further analyze the differences by the nature of attention in Appendix Section B.

By Gender Figure 3a reveals a striking deviation from educators' general tendency to prioritize lower-achieving students. In mixed-gender pairs, lower-achieving female students not only fail to receive additional attention, but they receive 1.4 p.p. *less* attention than their higher-achieving male peers ($t(1279) = -2.32$, $p = 0.02$). This pattern stands in stark contrast to how educators respond to lower-achieving male students, who receive a substantial 9.1 p.p. increase in attention compared to their higher-achieving female peers ($t(1262) = 14.54$, $p < 0.001$). This increase is driven by all categories of attention including content-focused attention (rf. Appendix Section B).

By Race Figure 3b reveals another notable deviation from educators' typical pattern of prioritizing lower-achieving students. In mixed-race pairs, lower-achieving Black students do not receive the additional attention we would expect; they receive about the same amount of attention as their non-Black higher-achieving peer. This is a striking difference to when both students are Black, where the lower-achieving Black student receives 8.6 p.p. more attention than their higher-achieving Black peer ($t(256) = 7$, $p < 0.001$). This increase in attention is also driven by all categories of attention (rf. Appendix Section B).

By EL Status Figure 3c shows that in mixed EL pairs, the lower-achieving student receives more attention, following the expected pattern. However, in pairs where both students have EL status, this trend reverses: the lower-achieving EL student receives 4.6 p.p. *less* attention than their higher-achieving EL peers ($t(740) = -5.56$, $p < 0.001$). This stands in contrast to all other pairings, where lower-achieving students receive more attention. The higher-achieving EL peer receives primarily more content-focused and relationship-building attention (rf. Appendix Section B).

Discussion

Educator attention is one of the most powerful resources for improving student learning and well-being. Yet, educator attention is often subtle and difficult to study. In the past, researchers needed to manually observe and record observations to identify patterns with potentially meaningful




Figure 3: Distribution of individual attention by achievement across student pairings. Error bars indicate standard errors. (a) Gender pairings: In mixed-gender groups, lower-achieving male students receive significantly more attention than their female counterparts. (b) Race pairings: Lower-achieving Black students in Black-only groups receive significantly more attention, but this pattern does not hold in mixed-race pairings. (c) English Learner (EL) status pairings: Lower-achieving EL students receive less attention when paired with another EL student. The lower-achieving non-EL student receives more attention in all settings.

consequences for students. The time costs of these observations often resulted in small samples, covering few educators and short time spans. The expansion of both large-scale virtual tutoring and computational methods for analyzing this data provides an unprecedented opportunity to understand patterns in educator-student interactions. Our study leverages these data and methods to complement existing observational insights, and provide new nuanced perspectives on educator attention.

Our findings reinforce some prior qualitative observations, such as gender-based disparities in educator attention (Beaman et al., 2006), and examine racial and EL disparities, areas that are understudied in large-scale settings. Our work shows that educators tend to allocate more attention to lower-achieving students, consistent with the benefits of need-based attention. However, we also identify notable deviations from this pattern. Lower-achieving female students in mixed-gender pairs receive significantly less attention than their higher-achieving male peers, while lower-achieving male students receive substantially more attention than their higher-achieving female

peers. A similar deviation emerges for lower-achieving Black students who do not receive additional attention when paired with a non-Black student, but do receive it when both students are Black. By contrast, higher-achieving EL students receive much more attention than their lower-achieving EL peers; one potential explanation for this pattern is that the tutor may find it easier to engage with the EL student who has the higher literacy level in these groups. By examining these disparities in a large-scale setting, our study extends prior research and sheds light on the interaction of multiple factors, including student needs and demographics, that shape educator-student interactions.

Several limitations must be acknowledged. Our study focuses on educator utterances, excluding student talk, which limits our ability to examine how student-initiated behaviors shape educator responses. Because educator-student interactions are inherently bidirectional (Brophy and Good, 1970), future work would benefit from incorporating student talk. For example, the gender-based differences in educator attention from Study 2 and 3 may be partially explained by differences in student behaviors, such as boys exhibiting more attention-seeking behaviors (Brophy and Good, 1970, Dobbs et al., 2004, Irvine, 1986). Additionally, the absence of educator demographic information prevents us from studying how educator identities influence patterns in interaction—an important factor in instructional effectiveness (Loeb et al., 2014).

Despite these limitations, our study demonstrates the power of large-scale data and computational methods to uncover subtle but consequential patterns in educational interactions. By identifying disparities in educator attention, this work surfaces empirical evidence that informs efforts to improve teaching practices and create more equitable learning environments.

Materials and Methods

Data and Processing. Our data is sourced from a U.S.-based early literacy, virtual tutoring provider that offers end-to-end services for school districts, including the tutoring platform, instructional materials, and tutors. While the tutoring provider offers different forms of tutoring, our study uses the 2-on-1 tutoring sessions where a tutor conducts tutoring with two students. These students are in K-2 grades and were identified as being eligible for tutoring because they were performing below the early literacy benchmark. The tutors are U.S. based, and are former classroom teachers and part-time teachers among others. The tutoring interactions are video-based, integrated on the providers' online platform.

Our study uses the tutor's audio recordings from the tutoring session, which were already individually saved from the students' audio recordings. We excluded tutoring sessions where the students could not be matched to their metadata information, such as their gender or race. Tutors could opt-in to providing their background information including demographics, however not all did. Therefore, we do not use the tutors' information in our analysis. We transcribed the tutor's audio with an automatic speech recognition model (Bain et al., 2023), where each transcribed utterance came with associated with start and end timestamps. We manually checked the quality of the transcripts to validate their useability.

The transcripts were de-identified by the tutoring provider, with their roster of student names. Because students are uniquely identifiable, their names are replaced with "[Student A]" and "[Student B]" in the transcripts. We automatically trimmed transcripts to when both students are present in the session based on the entry timestamps of the students because our study is interesting in how the tutor attends to different students in the session. Finally, we excluded sessions with durations that were less than half the planned tutoring session length because these sessions did not complete the desired amount of instruction. The final dataset size is 5,249 transcripts.

Development of Attention Framework Our attention framework was developed through an empirically grounded approach, drawing on both prior research and manual observations. To determine the recipient of attention, we built on existing work in referee classification (Jovanovic and op den Akker, 2004, Ouchi and Tsuboi, 2016), categorizing educator utterances based on the individual student. Observing that many educator utterances engaged both students simultaneously, we explicitly included a “both” category. For the nature of attention, we leveraged prior work on educator attention (Gunn et al., 2021, among others) and the training materials of the tutoring provider, both of which emphasize content instruction and relationship-building skills as central to interactions. Additionally, upon manual observation, we noted a substantial number of educator utterances focused on session management, aligning with educational theory on instructional practice and classroom management literature (Darling-Hammond and Bransford, 2007, among others). These observations informed our framework, ensuring it captured key theoretical and empirical dimensions of educator-student interactions.

Human Annotation of Recipient of Attention. A subset of 1,200 utterances were sampled from the corpus for annotation. Utterances were sampled with the constraint that at least 10 prior utterances as preceding context. We recruited 4 annotators, including two of the authors, who were trained on our attention framework and familiar with the data context to annotate for both the recipient and nature of attention. An additional category of “NA” if annotators were unsure about how to classify the recipient of the utterance. An initial round of 50 examples was annotated to assess inter-rater agreement, yielding Fleiss $\kappa = 0.66$ averaged over all categories. While this indicated substantial agreement (Landis and Koch, 1977), we had some cases marked with “NA” due to missing contextual cues (e.g., tone of voice) available in the original audio recordings. Therefore, we decided to doubly annotate all 1,200 examples and resolve disagreement manually. Before adjudication, the final agreement score was Cohen’s $\kappa = 0.613$. The two annotation authors resolved examples with disagreements and marked with “NA” by listening to the original audio recordings.

Human Annotation of Nature of Attention. The same subset of 1,200 utterances was used to annotate the nature of attention. We employed a Human+LLM annotation approach, leveraging recent advances in combining human expertise with AI to scale annotation efforts (Wang et al., 2024b, among others). We first prompted an LLM to annotate the dataset using our taxonomy, after which two co-authors reviewed and corrected the labels. Human annotators played a particularly important role to ensure that the final labels captured the nature of educator attention in the utterances.

Computational Annotation of Utterances To automatically label tutor utterances with the recipient and nature categories at scale, we split our annotated dataset into a 7:1:2 train:validation:test set and fine-tuned a RoBERTa-large model on the training examples (Liu et al., 2019). We created new token separators to separate the pretext context (10 preceding lines) and target utterance in the input sequences, formatting them as “[PRETEXT_TOKEN] {pretext} [TARGET] {target}” to match what annotators saw during labeling. We used a class-balanced loss function (Cui et al., 2019) during fine-tuning, as some categories appeared more frequently than others. We performed a hyper-parameter search over the class-balanced loss annealing term β and learning rate α , fine-tuning for 20 epochs under each configuration. The best model was selected based on its validation F1 score which was trained with $\beta = 0.9$ and $\alpha = 5e6$. This final SoS classification model achieved a test F1 score of 71.8, and test accuracy of 72.5. We applied this model to automatically label all tutor utterances in our larger corpus of 1,157,970 utterances for the downstream analyses in this work.

We compare the performance of our fine-tuned classifier to other modeling alternatives, including heuristic-based methods (e.g., detecting the presence of a student’s name) and k-shot prompting with large language models like Claude and GPT-4, and find that our classifier outperforms most methods. Additionally, our classifier offers a more cost-effective solution while maintaining high accuracy. For further details on our model evaluations, please see Appendix Section C.

Acknowledgments

We are grateful to our partners in this research, On Your Mark and Uplift Education. This research was graciously supported by the Overdeck Family Foundation, Smith Richardson Foundation, and Arnold Ventures. We also thank Lena Phalen and Professor Jennifer Osuna for insightful conversations.

References

- Joseph Allen, Anne Gregory, Amori Mikami, Janetta Lun, Bridget Hamre, and Robert Pianta. Observations of effective teacher–student interactions in secondary school classrooms: Predicting student achievement with the classroom assessment scoring system—secondary. *School psychology review*, 42(1):76–98, 2013.
- Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. Whisperx: Time-accurate speech transcription of long-form audio. *INTERSPEECH 2023*, 2023.
- Robyn Beaman, Kevin Wheldall, and Coral Kemp. Differential teacher attention to boys and girls in the classroom. *Educational review*, 58(3):339–366, 2006.
- Sondra H Birch and Gary W Ladd. The teacher-child relationship and children’s early school adjustment. *Journal of school psychology*, 35(1):61–79, 1997.
- Jere E Brophy and Thomas L Good. Teachers’ communication of differential expectations for children’s classroom performance: Some behavioral data. *Journal of educational psychology*, 61(5): 365, 1970.
- Joshua L Brown, Stephanie M Jones, Maria D LaRusso, and J Lawrence Aber. Improving classroom quality: Teacher influences and experimental impacts of the 4rs program. *Journal of educational psychology*, 102(1):153, 2010.
- Margaret Burchinal, Carollee Howes, Robert Pianta, Donna Bryant, Diane Early, Richard Clifford, and Oscar Barbarin. Predicting child outcomes at the end of kindergarten from the quality of pre-kindergarten teacher–child interactions and instruction. *Applied Development Science*, 12(3): 140–153, 2008.
- Raj Chetty, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter. Race and economic opportunity in the united states: An intergenerational perspective. *The Quarterly Journal of Economics*, 135(2): 711–783, 2020.
- Erica Cooper and Mike Allen. A meta-analytic examination of the impact of student race on classroom interaction. *Communication Research Reports*, 15(2):151–161, 1998.
- Kalena Cortes, Karen Kortecamp, Susanna Loeb, and Carly Robinson. A scalable approach to high-impact tutoring for young readers: Results of a randomized controlled trial. Technical report, National Bureau of Economic Research, 2024.

- Kalena E Cortes, Karen Kortecamp, Susanna Loeb, and Carly D Robinson. A scalable approach to high-impact tutoring for young readers. *Learning and Instruction*, 95:102021, 2025.
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019.
- Linda Darling-Hammond and John Bransford. *Preparing teachers for a changing world: What teachers should learn and be able to do*. John Wiley & Sons, 2007.
- Dorottya Demszky, Jing Liu, Heather C Hill, Dan Jurafsky, and Chris Piech. Can automated feedback improve teachers’ uptake of student ideas? evidence from a randomized controlled trial in a large-scale online course. *Educational Evaluation and Policy Analysis*, 46(3):483–505, 2024.
- Jennifer Dobbs, David H Arnold*, and Greta L Doctoroff. Attention in the preschool classroom: The relationships among child gender, child misbehavior, and types of teacher attention. *Early Child Development and Care*, 174(3):281–295, 2004.
- Carol S Dweck, William Davidson, Sharon Nelson, and Bradley Enna. Sex differences in learned helplessness: Ii. the contingencies of evaluative feedback in the classroom and iii. an experimental analysis. *Developmental psychology*, 14(3):268, 1978.
- National Center for Educational Statistics, United States. Office of Educational Research, Improvement. Center for Education Statistics, and Institute of Education Sciences (US). *Report on the condition of education*. US Department of Education, Office of Educational Research and Improvement . . . , 2024.
- FutureEd. How local educators plan to spend billions in federal covid aid, 2022. URL <https://www.future-ed.org/local-covid-relief-spending/>.
- Walter S Gilliam, Angela N Maupin, Chin R Reyes, Maria Accavitti, and Frederick Shic. Do early educators’ implicit biases regarding sex and race relate to behavior expectations and recommendations of preschool expulsions and suspensions. *Yale University Child Study Center*, 9(28):1–16, 2016.
- Roland H Good III, Ruth A Kaminski, Deborah Simmons, and Edward J Kame’enui. Using dynamic indicators of basic early literacy skills (dibels) in an outcomes-driven model: Steps to reading outcomes. *OSSC Bulletin*, 44(1):n1, 2001.
- Lucas Gortazar, Claudia Hupkau, and Antonio Roldán-Monés. Online tutoring works: Experimental evidence from a program with vulnerable children. *Journal of Public Economics*, 232:105082, 2024.
- Leiah Groom-Thomas, Chung Leung, Susanna Loeb, Cynthia Pollard, Nancy Waymack, and Sarah White. Challenges and solutions: Scaling tutoring programs. *Inter-American Development Bank*, 2023.
- Barbara Gunn, Keith Smolkowski, Lisa A Strycker, and Caroline Dennis. Measuring explicit instruction using classroom observations of student–teacher interactions (costi). *Perspectives on Behavior Science*, 44(2):267–283, 2021.
- Bridget K Hamre and Robert C Pianta. Can instructional and emotional support in the first-grade classroom make a difference for children at risk of school failure? *Child development*, 76(5):949–967, 2005.

- Kirby A Heller and Jacquelynne Eccles Parsons. Sex differences in teachers' evaluative feedback and students' expectancies for success in mathematics. *Child Development*, pages 1015–1019, 1981.
- Jacqueline J Irvine. Teacher–student interactions: Effects of student race, sex, and grade level. *Journal of Educational Psychology*, 78(1):14, 1986.
- Jennifer Jacobs, Karla Scornavacco, Charis Harty, Abhijit Suresh, Vivian Lai, and Tamara Sumner. Promoting rich discussions in mathematics classrooms: Using personalized, automated feedback to support reflection and instructional change. *Teaching and Teacher Education*, 112:103631, 2022.
- Susanne M Jones and Kathryn Dindia. A meta-analytic perspective on sex equity in the classroom. *Review of educational research*, 74(4):443–471, 2004.
- Natasa Jovanovic and Rieks op den Akker. Towards automatic addressee identification in multi-party dialogues. In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*, pages 89–92, 2004.
- Chen-Lin C Kulik and James A Kulik. Effects of ability grouping on secondary school students: A meta-analysis of evaluation findings. *American educational research journal*, 19(3):415–428, 1982.
- J. Richard Landis and Gary G. Koch. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174, 1977. ISSN 0006341X, 15410420. URL <http://www.jstor.org/stable/2529310>.
- James L Lewis, Robert K Ream, Kathleen M Bocian, Richard A Cardullo, Kimberly A Hammond, and Lisa A Fast. Con cariño: Teacher caring, math self-efficacy, and math achievement among hispanic english learners. *Teachers College Record*, 114(7):1–42, 2012.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019. URL <https://arxiv.org/abs/1907.11692>.
- Susanna Loeb, James Soland, and Lindsay Fox. Is a good teacher a good teacher for all? comparing value-added of teachers with their english learners and non-english learners. *Educational Evaluation and Policy Analysis*, 36(4):457–475, 2014.
- Clark McKown and Rhona S Weinstein. Teacher expectations, classroom context, and the achievement gap. *Journal of school psychology*, 46(3):235–261, 2008.
- Debra Myhill. Bad boys and good girls? patterns of interaction and response in whole class teaching. *British Educational Research Journal*, 28(3):339–352, 2002.
- Andre Nickow, Philip Oreopoulos, and Vincent Quan. The promise of tutoring for prek–12 learning: A systematic review and meta-analysis of the experimental evidence. *American Educational Research Journal*, 61(1):74–107, 2024.
- Hiroki Ouchi and Yuta Tsuboi. Addressee and response selection for multi-party conversation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2133–2143, 2016.
- RC Pianta. Classroom assessment scoring system™: Manual k-3. *Paul H Brookes Publishing*, 2008.
- Daniel L Reinholz and Niral Shah. Equity analytics: A methodological approach for quantifying participation patterns in mathematics classroom discourse. *Journal for Research in Mathematics Education*, 49(2):140–177, 2018.

- Carly D Robinson, Cynthia Pollard, Sarah Novicoff, Sara White, and Susanna Loeb. The effects of virtual tutoring on young readers: Results from a randomized controlled trial. *Educational Evaluation and Policy Analysis*, page 01623737241288845, 2024.
- James Soland and Lia E Sandilos. English language learners, self-efficacy, and the achievement gap: Understanding the relationship between academic and social-emotional growth. *Journal of Education for Students Placed at Risk (JESPAR)*, 26(1):20–44, 2021.
- Jim Soland. English language learners, self-efficacy, and the achievement gap. *Collaborative for student growth*, NWEA, 2019.
- Guillermo Solano-Flores, Maria Araceli Ruiz-Primo, Min Li, Xueyu Zhao, Chelsey Shade, and Ashley Chrzanowski. How equally do teachers distribute their attention across students classified as english learners (els) and their non-el peers in science classrooms? a frequency analysis of monolingual and bilingual teachers' interactions with different student grouping configurations. *International Multilingual Research Journal*, pages 1–15, 2024.
- Abhijit Suresh, Jennifer Jacobs, Charis Harty, Margaret Perkoff, James H Martin, and Tamara Sumner. The talkmoves dataset: K-12 mathematics lesson transcripts annotated for teacher and student discursive moves. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4654–4662, 2022.
- Rob Voigt, Nicholas P Camp, Vinodkumar Prabhakaran, William L Hamilton, Rebecca C Hetey, Camilla M Griffiths, David Jurgens, Dan Jurafsky, and Jennifer L Eberhardt. Language from police body camera footage shows racial disparities in officer respect. *Proceedings of the National Academy of Sciences*, 114(25):6521–6526, 2017.
- Rose E Wang, Ana T Ribeiro, Carly D Robinson, Susanna Loeb, and Dora Demszky. Tutor copilot: A human-ai approach for scaling real-time expertise. *arXiv preprint arXiv:2410.03017*, 2024a.
- Xinru Wang, Hannah Kim, Sajjadur Rahman, Kushan Mitra, and Zhengjie Miao. Human-llm collaborative annotation through effective verification of llm labels. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–21, 2024b.
- Carol Watson, Jenna Plump, and James Durham. Gender bias in the middle level classroom: The intersection of observation, teacher self-perceptions, and student perceptions. *Journal of Research in Education*, 31(1):124–150, 2022.
- Michael Younger, Molly Warrington, and Jacquetta Williams. The gender gap and classroom interactions: reality and rhetoric? *British journal of Sociology of Education*, 20(3):325–341, 1999.

APPENDIX

A Data

A.1 Context

Our data is sourced from a U.S.-based early literacy, virtual tutoring provider that offers end-to-end services for school districts, including the tutoring platform, instructional materials, and tutors. While the tutoring provider offers different forms of tutoring, our study uses the 2-on-1 tutoring sessions where a tutor conducts tutoring with two students. Each tutoring session is scheduled to be 20 minutes and the tutoring provider prioritizes consistent tutor assignments for students. The tutoring sessions take place from the latter half of the academic year. The research team executed Data Use Agreements with the tutoring provider and a charter school network that outlined the allowable usage of the data to improve instruction in collaboration with an educational agency. Following FERPA guidelines, we were eligible to engage in secondary data analysis with student data, which is what we did for this study. The tutoring curriculum includes a focus on phonics, phonological awareness, and fluency. These students are in K-2 grades and were identified as being eligible for tutoring because they were performing below the early literacy benchmark. The tutors are U.S. based, and are former classroom teachers and part-time teachers among others; they receive training and professional development, focused on content knowledge, building relationships, and effective delivery of the intervention. The tutoring interactions are video-based, integrated on the providers' online platform. Additional demographic information on students and tutors can be found in Table A.1.

A.2 Transcription and Metadata

The tutor's audio from 2:1 sessions was transcribed with WhisperX (Bain et al., 2023), associated with timestamps. The students' audio suffered from low audio quality and was not transcribed: tutoring happened in school, therefore the loud background noise made transcription unreliable.

Each student is linked with metadata on their demographic information, which includes information about the student's gender, race, English language learner (ELL) status, socio-economic status, and special education status. Each student is also linked with their mid-of-year (MOY) and end-of-year (EOY) Dynamic Indicators of Basic Literacy Skills (DIBELS) performance, which is a measure of the student's literacy skills. The MOY DIBELS examination was taken prior to the tutoring sessions we use for this study; this acts as a control on the student's baseline performance, which may impact how much attention they receive from the tutor. We excluded tutoring sessions where the students could not be matched to their metadata information. Tutors could opt-in to providing their background information including demographics, however not all did. Therefore, we do not use the tutors' information in our analysis. Table A.2 reports descriptive statistics on sessions and transcripts.

A.3 Preprocessing

The transcripts were de-identified by the tutoring provider, with their roster of tutor and student names. Because students are uniquely identifiable, their names are replaced with "[Student A]" and "[Student B]" in the transcripts. Because we're interested in how the tutor talks with more than one student in the session, we trimmed transcripts to when both students are present in the session; we automatically do this using the entry timestamps of the students. Finally, we excluded sessions with durations that were less than half the planned session length (10 minutes). The final dataset size is ~ 5.2k transcripts.

Tutors		Students	
# of Tutors	125	# of Students	757
Demographics		Grade	
Asian American	6%	Kindergarten	207 (27%)
Black	35%	First	293 (39%)
Latina/o/x	12%	Second	257 (34%)
Multiracial	8%	Demographics	
White	45%	Female	53%
Educational Status		Asian American	1%
Grad Studying	8%	Black	22%
College Grad	47%	Latina/o/x	70%
College Studying	13%	Multiracial	3%
High School Grad	6%	White	4%
Teaching Experience		Multilingual Learners	35%
Former Teacher	33%	Student with Disabilities	5%
Part-time Teacher	19%	LIEM	89%
Other Education Professional	21%		

Table A.1: **Statistics on identifiable teachers and students.** Tutor demographics are reported for those who could be matched to transcripts using the platform ID crosswalk ($n = 89$ out of 125). For each demographic, we use the naming convention from (Robinson et al., 2024). LIEM = Low-income economically marginalized.

Measure	Count	Mean	Standard Deviation
Session Duration (in seconds)	5,604,605	1067.75	180.68
Number of Words	6,723,347	1280.88	442.89
Number of Utterances	1,157,970	220.61	71.74

Table A.2: **Descriptive statistics of transcripts.**

A.4 Lower- vs. Higher-Achieving Students

In our paired tutoring context, we define the lower-achieving student as the one with the lower baseline achievement score. This baseline is determined by the student’s literacy test score, specifically the DIBELS composite score, taken prior to data collection (Good III et al., 2001).

Table A.3 shows the statistics of the baseline achievements for the lower-achieving student and the higher-achieving student in standard deviation units. The DIBELS composite scores are standardized within each grade since the DIBELS composite score, which combines multiple subtests (e.g., phonemic awareness, oral reading fluency) varies by expected skill level per grade.

Achievement Status	Mean	Standard Deviation
Higher	0.48	0.95
Lower	-0.35	0.81

Table A.3: Descriptive statistics of the lower- vs. higher-achieving student in each tutoring pair.

B Additional Attention Results

B.1 Nature of Attention from Study 2

The main study reports the total amount of attention students received by their demographic background and pair composition receive. Here, we report the amount of attention broken down to our three main categories of attention: Content, relationship-building, and management. We use the same model as in Study 2 of our paper, included here for reference: We model the amount of that attention type, y_{ijk} , a student i receives when partnered with student j in session k as a function of: an indicator of their pair composition I_{ijk} (e.g., female in same-gender or mixed-gender pairings), their own baseline achievement D_i , their partner’s achievement D_j and their relative achievement (lower or higher) R_{ij} . We cluster standard errors at the pair level, e_{ij} .

$$y_{ijk} = \beta_1 I_{ijk} + \beta_2 D_i + \beta_3 D_j + \beta_4 R_{ij} + e_{ij}$$

Figure B.1 reports the estimated differences in educator attention based on gender. We find that female students in mixed-gender pairings receive much less content-focused and management-focused attention, than male students.

Figure B.2 reports based on race. We find that pairs of non-Black students receive much less individual attention on content.

Figure B.3 reports based on the students’ EL status. We find that pairs of non-EL students receive much less individual attention on content as well.

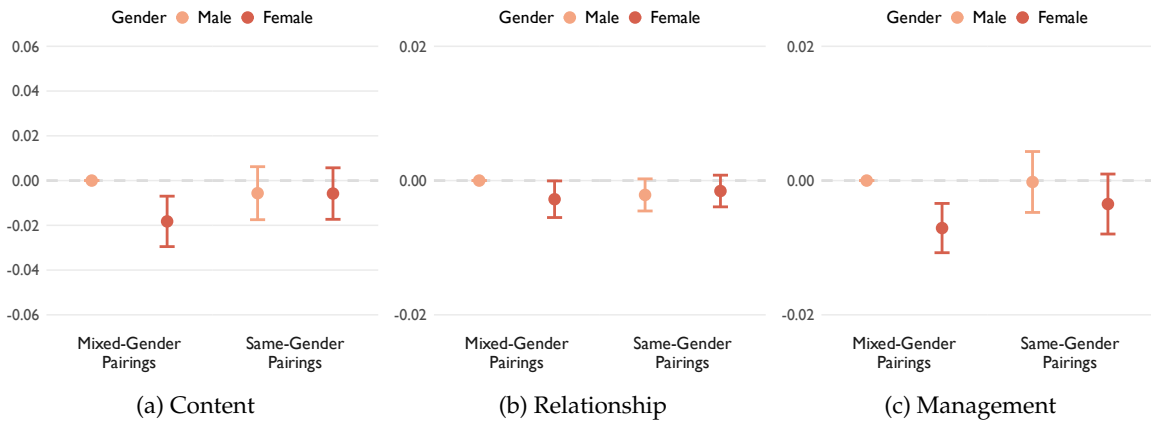


Figure B.1: Nature of attention by gender.




Figure B.2: Nature of attention by race.




Figure B.3: Nature of attention by EL status.

B.2 Nature of Attention from Study 3

The main study reports aggregate educator attention based on student demographics and pair composition. Here, we break down these findings by the nature of attention (content, relationship-building, and management) reported in Figure B.4.

By Gender The main study reports that in mixed-gender pairs, lower-achieving female students not only fail to receive additional attention but receive significantly *less* attention than their higher-achieving male peers. Figure B.4c reveals that this disparity is primarily driven by an increase in management-focused attention directed toward the male student. In contrast, lower-achieving male students receive substantially more attention than their higher-achieving female peers, and this increase is observed across all attention categories.

By Race The main study reports that in mixed-race pairs, lower-achieving Black students do not receive the additional attention we would expect. They receive about the same amount of attention as their non-Black higher-achieving peer. However, when both students are Black, the lower-achieving Black student receives more attention than their higher-achieving Black peer, and this is driven by all forms of attention.

By EL Status The main study reports that in pairs with mixed EL status, the lower-achieving EL student receives more attention. This is primarily driven by attention on content and relationship-building. However, in pairs where both students have EL status, the trend of attending to the lower-achieving student reverses: the lower-achieving EL student receives *less* attention than their higher-achieving EL peers, especially in content and relationship-building.

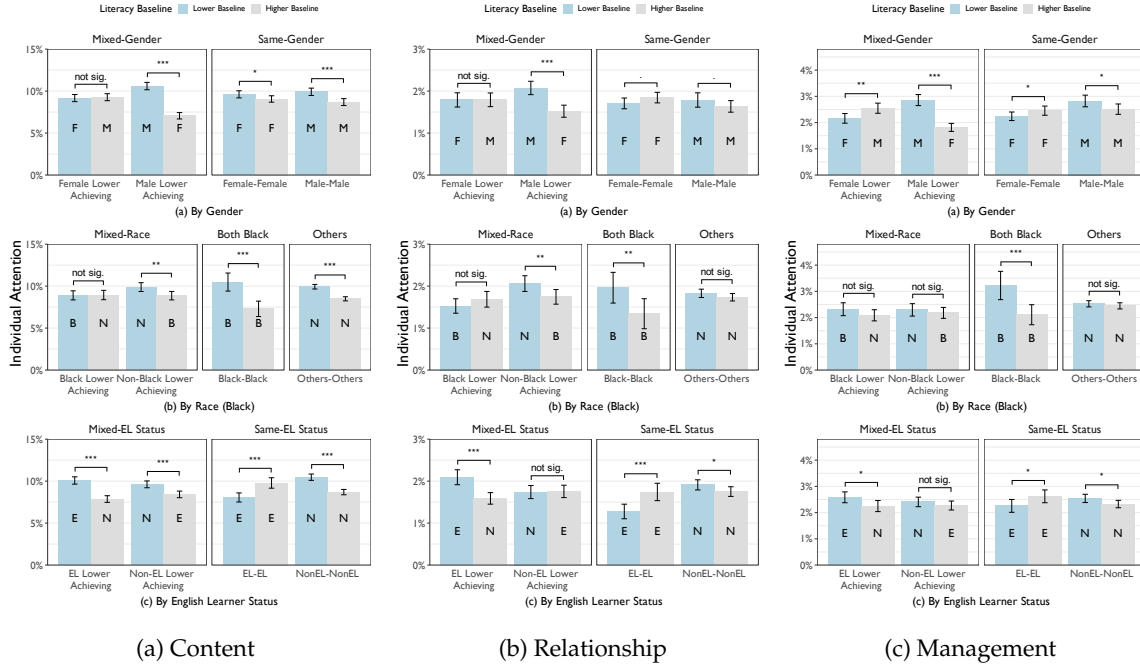


Figure B.4: Nature of attention by student achievement relative to their peer.

B.3 Results on Attention for “Both Students” and “One of the Students”

Our framework additionally identifies when educator utterance is directed to both students or directed to an ambiguous recipient. We report those results in this section.

Pairings	Mean	Standard Deviation	n
Mixed-Gender	0.369	0.160	2543
Female-Female	0.378	0.157	1606
Male-Male	0.390	0.157	1100
Mixed-Race	0.374	0.157	1619
Black-Black	0.350	0.146	257
Others-Others	0.379	0.160	3373
Mixed-EL Status	0.383	0.158	2363
EL-EL	0.403	0.173	741
NonEL-NonEL	0.359	0.151	2101

Table B.4: Descriptive statistics on “both students” as the attention recipient.

Category	Pairings	Estimate
By Gender	Female-Female	-0.015*** (0.004)
	Male-Male	-0.042*** (0.005)
By Race (Black)	Black-Black	0.034*** (0.009)
	Others-Others	-0.021*** (0.004)
By EL Status	EL-EL	0.011** (0.006)
	NonEL-NonEL	0.001 (0.004)
Observations		5,205
R ²		0.025
Adjusted R ²		0.024
Residual Std. Error		0.129 (<i>df</i> = 5198)
F Statistic		22.443*** (<i>df</i> = 6; 5198)
<i>Note:</i>	* <i>p</i> < 0.1; ** <i>p</i> < 0.05; *** <i>p</i> < 0.01	

Table B.5: **Regression-based check on ambiguous “Only one” recipient category.** The heterogeneous pairing (e.g. mixed-gender) is set as the baseline.

Table B.4 reports the amount of educator attention directed to both students across different pair compositions.

The ambiguous “One of the Students” category captures utterances directed to only one student but the exact recipient is unclear. One potential concern is whether the amount of ambiguous references is large enough to dispute our earlier findings. We conduct a regression-based check to test the robustness of our results. We model of the amount of ambiguous attention, y_k in session k as function of indicators of their Gender pair composition G_k (i.e. both female, both male, mixed-gender), Race (black) pair composition R_k , and EL status pair composition E_k .

$$y_k = \beta_1 G_k + \beta_2 R_k + \beta_3 E_k + e_k$$

Table B.5 reports the coefficients of the regression. We compare the coefficients against Study 2’s results and confirm that our findings are robust. Take gender pairings as an example. Table B.5 shows that compared to both-male sessions, mixed-gender sessions have 4.2 p.p. more utterances classified as ambiguous. Meanwhile, we find in Study 2 that in mixed-gender pairings female students receive 5.2 p.p. less attention than their male partners. Even if we assume that all of the additional 4.2 p.p. of ambiguous utterances were actually directed to female students (but misclassified), females would still receive 1 p.p. less attention than their male partners. Thus, our earlier findings still hold.

C Additional Model Evaluations

To determine the best classification approach, we evaluated various models, methods and heuristics.

We compare the following approaches in the table:

- Name in Text: A simple heuristic where the recipient is identified based solely on the presence of names in the target utterance text. If '[Student A]' appears in the utterance, then the utterance is classified as addressing Student A, same with '[Student B]' for Student B. If both names appear, then the utterance is classified as addressing both students. If no names appear, then the utterance is classified as addressing one of the students.
- Name in Text & Context: A revised heuristic that follows the same rules as the previous heuristic, but check for names in the target utterance text and the preceding context (10 previous conversation turns).
- RoBERTa (Liu et al., 2019): Fine-tuned on the train split of our annotated dataset.
- GPT-4, Claude-Haiku, Claude-Sonnet, Claude-Opus: Closed-source models with varying performance and cost implications. We tested these models with a zero-shot classification setup or k -shot setup. Prompts for these setups are shown in Figures C.5 (0-shot), C.6 (1-shot) and C.7 (3-shot).

We report the F1 score and the average API cost per 100 transcripts, based on pricing from OpenAI and Anthropic. The performance metrics are reported in Table C.6 on our original dataset for classifying recipient of attention. Given the cost-effectiveness and performance, we chose to finetune RoBERTa models for our study since they provided similar performance to the best closed-source options without incurring additional API expenses.

Method	k	F1	Cost (\$)
Name in text	-	0.279	-
Name in text & context	-	0.435	-
RoBERTa	-	0.718	-
GPT4	0	0.577	62.1
GPT4	1	0.607	83.4
GPT4	3	0.654	128.0
Claude-Haiku	0	0.514	1.82
Claude-Haiku	1	0.443	2.37
Claude-Haiku	3	0.530	3.53
Claude-Sonnet	0	0.629	21.83
Claude-Sonnet	1	0.608	28.46
Claude-Sonnet	3	0.725	42.32
Claude-Opus	0	0.655	109.16
Claude-Opus	1	0.670	142.31
Claude-Opus	3	0.698	211.59

Table C.6: **Classification performance.** Cost is reported as estimated cost for annotating 100 transcripts with the method.

Prompt for zero-shot classification

Your task is to read the following conversation snippet and classify whom the tutor is talking to. The conversation comes from a K-2 early literacy tutoring session between a tutor and two students. These students have been de-identified as [Student A] and [Student B].

The possible labels are:

- 0: The tutor is addressing both students. e.g., "Let's do it together, [Student A], [Student B] and me."
- 1: The tutor is addressing Student A. e.g., "Okay, [Student A], it's your turn."
- 2: The tutor is addressing Student B. e.g., "Good job, [Student B]."
- 3: The tutor is addressing one of the students, but it is unclear which one. e.g., "Let's wait for him."

Only output the label number. Do not output anything else.

Context: {context}

Text: {text}

Label (number):

Figure C.5: **Prompt for zero-shot classification used with GPT-4, Claude-Haiku, Claude-Sonnet, Claude-Opus.** {context} is the placeholder for the educator utterances (10 preceding lines) leading up to the target utterance for classification. {text} is the educator utterance we aim to assign label.

Prompt for k -shot classification ($k = 1$)

Your task is to read the following conversation snippet and classify whom the tutor is talking to. The conversation comes from a K-2 early literacy tutoring session between a tutor and two students. These students have been de-identified as [Student A] and [Student B].

The possible labels are:

- 0: The tutor is addressing both students. e.g., "Let's do it together, [Student A], [Student B] and me."
- 1: The tutor is addressing Student A. e.g., "Okay, [Student A], it's your turn."
- 2: The tutor is addressing Student B. e.g., "Good job, [Student B]."
- 3: The tutor is addressing one of the students, but it is unclear which one. e.g., "Let's wait for him."

Only output the label number. Do not output anything else.

Context: Don't do it, don't do it, don't do it, don't do it, don't do it. She can circle it, but this is, it's for [Student B] to circle. It's for [Student B] to circle. It's hers, because Gamela had a turn. We got to learn to take turns.

Text: You got it, you have it, [Student B].

Label (number): 2

Context: {context}

Text: {text}

Label (number):

Figure C.6: **Prompt for k -shot classification ($k = 1$) used with GPT-4, Claude-Haiku, Claude-Sonnet, Claude-Opus.** {context} is the placeholder for the educator utterances (10 preceding lines) leading up to the target utterance for classification. {text} is the educator utterance we aim to assign label.

Prompt for k -shot classification ($k = 3$)

Your task is to read the following conversation snippet and classify whom the tutor is talking to. The conversation comes from a K-2 early literacy tutoring session between a tutor and two students. These students have been de-identified as [Student A] and [Student B].

The possible labels are:

- 0: The tutor is addressing both students. e.g., "Let's do it together, [Student A], [Student B] and me."
- 1: The tutor is addressing Student A. e.g., "Okay, [Student A], it's your turn."
- 2: The tutor is addressing Student B. e.g., "Good job, [Student B]."
- 3: The tutor is addressing one of the students, but it is unclear which one. e.g., "Let's wait for him."

Only output the label number. Do not output anything else.

Context: Don't do it, don't do it, don't do it, don't do it, don't do it. She can circle it, but this is, it's for [Student B] to circle. It's for [Student B] to circle. It's hers, because Gamela had a turn. We got to learn to take turns.

Text: You got it, you have it, [Student B].

Label (number): 2

Context: Who is like this or how? [Student A] has hair. How does she have hair? [Student A] has hair. have, however, horses have We missed it too. We can't- it has to start with the letter H. I'm gonna put headphones. This is funny. Oh [Student B], why do I keep doing that? There you go.

Text: So that's- that's- that's the sentence.

Label (number): 0

Context: This just helps us kind of map out the sounds that we hear. Oh, your word is bonnet. I'm going to move them for you. No, your word is kitten, [Student B]. Are you missing any? Good job, [Student B]. OK. [Student A], I'm going to tell you your word one more time. Bonnet. So let's see where we can fix it.

Text: Because you put bonnet.

Label (number): 1

Context: {context}

Text: {text}

Label (number):

Figure C.7: **Prompt for k -shot classification ($k = 3$) used with GPT-4, Claude-Haiku, Claude-Sonnet, Claude-Opus.** {context} is the placeholder for the educator utterances (10 preceding lines) leading up to the target utterance for classification. {text} is the educator utterance we aim to assign label.