



The NCTE Transcripts: A Dataset of Elementary Math Classroom Transcripts

Dorottya Demszky
Stanford University

Heather Hill
Harvard University

Classroom discourse is a core medium of instruction --- analyzing it can provide a window into teaching and learning as well as driving the development of new tools for improving instruction. We introduce the largest dataset of mathematics classroom transcripts available to researchers, and demonstrate how this data can help improve instruction. The dataset consists of 1,660 45-60 minute long 4th and 5th grade elementary mathematics observations collected by the National Center for Teacher Effectiveness (NCTE) between 2010-2013. The anonymized transcripts represent data from 317 teachers across 4 school districts that serve largely historically marginalized students. The transcripts come with rich metadata, including turn-level annotations for dialogic discourse moves, classroom observation scores, demographic information, survey responses and student test scores. We demonstrate that our natural language processing model, trained on our turn-level annotations, can learn to identify dialogic discourse moves and these moves are correlated with better classroom observation scores and learning outcomes. This dataset opens up several possibilities for researchers, educators and policymakers to learn about and improve K-12 instruction.

VERSION: November 2022

Suggested citation: Demszky, Dorottya, and Heather C. Hill. (2022). The NCTE Transcripts: A Dataset of Elementary Math Classroom Transcripts. (EdWorkingPaper: 22-682). Retrieved from Annenberg Institute at Brown University: <https://doi.org/10.26300/npxh-kf69>

The NCTE Transcripts: A Dataset of Elementary Math Classroom Transcripts

Dorottya Demszky
Stanford University
ddemszky@stanford.edu

Heather Hill
Harvard University
heather.hill@gse.harvard.edu

Abstract

Classroom discourse is a core medium of instruction – analyzing it can provide a window into teaching and learning as well as driving the development of new tools for improving instruction. We introduce the largest dataset of mathematics classroom transcripts available to researchers, and demonstrate how this data can help improve instruction. The dataset consists of 1,660 45-60 minute long 4th and 5th grade elementary mathematics observations collected by the National Center for Teacher Effectiveness (NCTE) between 2010-2013. The anonymized transcripts represent data from 317 teachers across 4 school districts that serve largely historically marginalized students. The transcripts come with rich metadata, including turn-level annotations for dialogic discourse moves, classroom observation scores, demographic information, survey responses and student test scores. We demonstrate that our natural language processing model, trained on our turn-level annotations, can learn to identify dialogic discourse moves and these moves are correlated with better classroom observation scores and learning outcomes. This dataset opens up several possibilities for researchers, educators and policymakers to learn about and improve K-12 instruction. The data and its terms of use can be accessed here: <https://github.com/ddemszky/classroom-transcript-analysis>

1 Introduction

Improving K-12 mathematics instruction in the aftermath of the Covid-12 pandemic is a major national priority, drawing support from both the U.S. government (U.S. Department of Education¹) and

¹<https://www.ed.gov/news/press-releases/us-department-education-announces-over-220-million-dollars-investments-government-private-and-public-sectors-support-student-recovery>

major foundations (e.g., Gates, Spencer²). A key step in this direction is to measure and facilitate the use of effective mathematics teaching practices, an effort that draws on a long history of research (e.g., Brophy, 1984; Sedova et al., 2019). Instructional measurement has traditionally relied on resource-intensive classroom observation. Recent natural language processing models, trained on manually scored classroom transcripts, enable measuring effective instructional practices in scalable and adaptable ways (Kelly et al., 2018; Suresh et al., 2019; Demszky et al., 2021b; Alic et al., 2022; Hunkins et al., 2022). However, a common barrier to evaluating such measures is the lack of comprehensive data sources that link classroom transcripts to external variables, such as student and teacher demographics and learning outcomes.

To address this, we introduce a dataset of 1,660 U.S. 4th and 5th grade elementary math classroom transcripts collected by the National Center for Teacher Effectiveness (NCTE) between 2010-2013 (Kane et al., 2015). The anonymized transcripts represent data from 317 teachers across 4 school districts serving largely historically marginalized student populations. The transcripts are associated with a wide range of metadata: (i) turn-level annotations for discourse moves, (ii) classroom observation scores, (iii) questionnaires that capture teacher background, beliefs and classroom practices, (iv) student administrative data, (v) questionnaires describing student background and classroom experiences, (vi) value added scores, which estimate teachers' contribution to students' academic performance. To our knowledge, this is the largest dataset of math classroom transcripts with linked outcomes available to researchers.

To illustrate how this data can be used to identify

²<https://www.spencer.org/news/announcing-covid-19-related-special-grant-cycle>

effective instructional practices, we build classifiers for discourse moves and validate these measures by correlating them with instructional outcomes. These discourse moves include on vs off task instruction, teachers' uptake of student ideas (Demszky et al., 2021b), teachers' focusing questions (Alic et al., 2022) and student reasoning — the latter three of which are indicators of dialogic instruction, where students are active participants of the learning process (Bakhtin, 1981; Nystrand et al., 1997; Wells, 1999; Alexander, 2008).

We show that a RoBERTa classifier can learn to predict these discourse moves with moderate to high accuracy, leaving some room for improvement for future work. Importantly, we find that predictions for all of these discourse moves correlate significantly with classroom observation scores that measure instructional quality, teacher sensitivity, and classroom climate, among other items, while controlling for teacher and classroom covariates. Predictions for dialogic moves (teacher uptake, focusing questions and student reasoning) also show a significant positive correlation with teachers' value added scores. Taken together, these results demonstrate the value of this dataset for developing measures that can help us understand and facilitate effective instruction, for example, by powering automated feedback tools for teachers (Suresh et al., 2021; Demszky et al., 2021a).

2 Related Work

We provide an overview of related corpora and methods pertaining to the computational analysis of classroom discourse.

2.1 Related Corpora

There is a wide range conversational datasets available to researchers, capturing phone conversations (Godfrey and Holliman, 1997), task oriented dialogue (Budzianowski et al., 2018), meeting transcripts (Braley and Murray, 2018), among many others. These datasets can provide valuable insights about social dynamics in conversations, but to understand teaching and learning, we need to capitalize on datasets from the educational domain.

Conversational datasets in the education domain are scarce, due to resource intensiveness of data collection and privacy protections. The Measures of Effective Teaching (MET) dataset (Kane et al., 2013) is the most similar to the NCTE data in terms of availability of outcomes. The MET data

contains 2,500 4-9th grade classroom recordings collected between 2009-2011 in six U.S. school districts. The recordings cover a variety of subjects, including English Language Arts and mathematics, and the recordings come with classroom observation scores, teacher and student demographic data and teacher value added scores. Although this dataset is a rich resource for studying instruction, it is challenging to work with, as it requires a paid subscription and interfacing with the data via a remote server that does not currently support many types of machine learning analyses. Furthermore, the MET data does not include transcripts; although subsets of the data have been transcribed as part of various research projects (Liu and Cohen, 2020; Hunkins et al., 2022), these transcripts are only accessible to these respective research teams.

The TalkMoves dataset (Suresh et al., 2022) is currently the largest publicly available collection of transcripts of U.S. math instruction. The TalkMoves data contains 567 K-12 math classroom transcripts, annotated for talk moves based on accountable talk theory (Michaels et al., 2010) and dialog acts from Switchboard (Jurafsky et al., 1997). The NCTE dataset complements the TalkMoves data in terms of availability of observation scores and outcomes. Other publicly available datasets of educational interactions outside of the K-12 domain include the CIMA tutoring dataset (Stasaski et al., 2020), STEM lectures captured in the DRYAD dataset (Reimer et al., 2016) and the Coursera forum discussion dataset (Rossi and Gnawali, 2014), among others.

2.2 Computational Analysis of Educational Interactions

Our detection of discourse moves relates to a long-standing line of work on dialog act classification. The dialog act classification literature focuses on domain-agnostic dialog acts, e.g. acknowledgment, repetition, questioning, etc. In this work, we focus on detecting discourse moves that are indicators of better math instruction, similarly to the TalkMoves project (Suresh et al., 2019). Suresh et al. (2019) train transformer-based classifiers, including BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), on annotations for six accountable talk moves, such as keeping everyone together, getting students to relate, revoicing, pressing for reasoning. The NCTE data comes with annotations for related discourse moves (e.g. revoicing in TalkMoves and

uptake in NCTE are similar constructs); predicting accountable talk moves on the NCTE data and studying correlations with outcomes is a promising direction for future work.

To better understand linguistic indicators of teacher effectiveness, [Liu and Cohen \(2020\)](#) analyze English Language Arts classes in the MET dataset ([Kane et al., 2013](#)) using topic models, LIWC ([Pennebaker et al., 2001](#)) and an open-ended question classifier. The authors conduct factor analysis on several linguistic features and find that a factor indicating interactive, student centered instruction correlates positively with teachers’ value added scores. [Hunkins et al. \(2022\)](#) annotate transcripts of 156 video clips from 6-8th grade math classrooms in the MET dataset for teacher talk moves that support belonging and inclusivity, such as praise, admonishment, controlling language and learning mindset supportive language. They build a random forest classifier to predict these talk moves, and find that admonishment, for example, has a negative correlation with students’ perception of the classroom environment. Finally, closely related work on the NCTE data has demonstrated the positive correlation between teacher uptake ([Demszky et al., 2021b](#)) and focusing questions ([Alic et al., 2022](#)) and instructional outcomes. In this work, we expand our correlational analyses to all discourse moves annotated on the NCTE dataset.

3 Dataset Description

The dataset consists of 1,660 anonymized transcripts of whole lessons, collected as part of the National Center for Teacher Effectiveness (NCTE) Main Study ([Kane et al., 2015](#)).³ The observations took place between 2010-2013 in 4th and 5th grade elementary math classrooms across four districts serving largely historically marginalized students. In the first two years, teachers and students were assigned to each classroom according to their school’s usual procedure for forming classes. In the third year, the NCTE project team randomly assigned teachers to rosters of students provided by the school.

Table 1 provides key statistics about the transcripts, as well as teacher and student demographics. The majority of teachers in the data are white (65%) and female (84%). Whereas the student

³Parents and teachers gave consent for the study (Harvard IRB #17768), and for de-identified data to be publicly shared for research.

Transcripts	
# of Transcripts	1660
Year 1	697
Year 2	616
Year 3	347
Avg # Transcripts Per Teacher	5.24 (±2.46)
Avg # of Turns	350 (±186)
Avg % of Teacher Turns	50.2% (±4.8%)
Avg # of Words	5733 (±1782)
Avg % of Teacher Words*	87.7% (±7%)
Teachers	
# of Teachers	317
% Male	16%
% Black	22%
% Asian	3%
% Hispanic/Latinx	3%
% White	65%
Avg # of Years of Experience	10.23 (7.28)
U.grad or Grad Degree in Math	6%
BA in Education	53%
Masters Degree	76%
Students	
# of Students	10,817
Grade	4th (51%) 5th (47%)
% Male	50%
% African American	43%
% Asian	8%
% Hispanic/Latinx	23%
% White	23%
% Free or Reduced Lunch	67%
% Special Education Status	13%
Limited English Proficiency	21%

Table 1: Statistics on transcripts and on teachers and students who are mappable to these transcripts via administrative data. For each demographic, we use the naming convention from [Kane et al. \(2015\)](#). Percentages for student grade levels do not add up to 100% due to missing values.

body is equally split in terms of gender, the majority are students of color (43% African American, 23% Hispanic/Latinx, 8% Asian) and receive free or reduced lunch (67%). This disparity between student-teacher demographics is in accordance with national statistics⁴ (Schaeffer, 2021).

3.1 Transcription & Anonymization

Lessons were captured by ThereNow using its Iris system⁵, which featured three cameras, a lapel microphone worn by teachers, and a bidirectional microphone for capturing student talk. The recordings were transcribed by professional transcribers working under contract to a commercial transcription company. Transcripts are fully anonymized: student and teacher names are replaced with terms like “Student J”, “Teacher” or “Mrs. H”. Inaudible talk, due to classroom noise and far field audio is transcribed as *[Inaudible]*. If the transcriber was unsure of a particular word, they transcribed it within brackets, e.g. *It is a city surrounded by [water]*. Square brackets are also used for other transcriber comments, such as *[crosstalk]*, and *[laughter]*. Almost all teacher talk and the majority of student talk could be transcribed: only 4% of teacher utterances and 21% of student utterances contain an *[Inaudible]* marker. Transcripts contain 5,733 words on average, 87.7% of which are spoken by the teacher.

3.2 Linked Variables

The transcript data comes with a uniquely rich source of linked variables, summarized in Table 2. These variables include turn-level annotations for various discourse features, classroom observation scores, demographic information about students and teachers, survey data and value-added scores. Please refer to Kane et al. (2015) for a full documentation of these variables, except for the turn-level annotations, which we describe below.

Turn-level annotations. Table 3 includes examples for each discourse feature, annotated at the turn level. In prior work, experts annotated a sample of 2,348 utterance pairs — exchanges between students and teachers — for on vs off task instruction, teachers’ uptake of student ideas (Demszky et al., 2021b) and focusing questions (Alic et al., 2022). The annotation process is described in its

⁴<https://nces.ed.gov/>

⁵ThereNow is no longer in business, but their technology is now used by IrisConnect: <https://www.irisconnect.com/uk/products-and-services/video-technology-for-teachers/>

respective papers. We include the coding scheme on our Github along with the dataset.

In a separate annotation process, experts coded 2,000 student utterances for student reasoning, a key Common Core aligned student practice. The coding process was based on the MQI classroom observation item “Student Provide Explanations” (Hill et al., 2008). To create a sample for labeling, we (i) hold out half of the transcripts for testing, (ii) from the remaining half, sample 30% of transcripts from the top quartile in terms of their Student Provide Explanations MQI score, and 70% from the rest, (iii) filter out student utterances shorter than 8 words, since they are unlikely to substantive reasoning, (iv) randomly sample up to 5 student utterances from each transcript, to balance representation across transcripts, (v) randomly sample 2,000 student utterances by assigning sampling weights proportionate to classroom diversity: combined percentage of African American and Hispanic/Latinx students in classroom. Each example was randomly assigned to one of two math coaches, who are also experts in the MQI coding instrument. Inter-rater agreement on calibration set (n=200) was 90%.

4 Computational Analysis

We illustrate how the NCTE transcript dataset can serve as a valuable resource for developing computational measures of classroom discourse. First, we train models that automatically identify dialogic discourse moves by leveraging the turn-level annotations. We then study how these discourse moves correlate with observation scores and teachers’ value added scores.

4.1 Pre-Processing Annotations

We binarize annotations for each discourse feature in order to make the classification task consistent across features for the purposes of this paper. We also found that providing teachers with feedback using only their positive examples can be effective (Demszky et al., 2021a), and thus in a first pilot experiment, one may not need to preserve fine-grained distinctions between negative, mediocre and positive examples.⁶

Labels for on vs off task instruction and student reasoning are binary, so we simply consider the majority rater label for these discourse moves. Labels

⁶Using these distinctions in teacher feedback is a promising direction of future work, given that contrasting examples can be an effective pedagogical tools (Schwartz et al., 2011; Sidney et al., 2015).

Variable	Description
Turn-level annotations	Annotations for on vs off task instruction, uptake of student contributions, focusing questions and student reasoning.
Transcript-level observation scores	Observation scores by expert raters using two instruments: CLASS (Pianta et al., 2008) and MQI (Hill et al., 2008).
Student questionnaires	Student survey responses about the classroom experience and their household.
Value-added scores	Teachers’ value added scores – i.e., an estimate of teachers’ contribution to students’ test performance.
Student administrative data	Administrative data on students, e.g. their test scores and demographic information.
Teacher questionnaires	Teacher’s self reported information about their background, beliefs and classroom practices.

Table 2: Variables linked to the NCTE transcript data. For full documentation, please refer to Kane et al. (2015).

Discourse Feature	Example
Student on Task	S: We both have the same number of blue, and red, and yellow.
Teacher on Task	T: Good, find the range. Find the range. Remember it’s the span of the least to the greatest number.
Student Reasoning	S: Because if you add ninety-eight hundredths and five hundredths, I think it’s going to add up to, like, it’s almost – it’s going to almost add up to a hundred.
High Uptake	S: I think these are Y axis and the X axis. T: They do. Sometimes they refer to them as X and Y axis. It depends on the type of graph. Okay. You ready?
Focusing Question	S: Four fifths – no, 80 percent. T: How come you can’t put it there?

Table 3: Examples for each discourse feature annotated at the turn-level. See Demszyk et al. (2021b) and Alic et al. (2022) for more details.

for uptake and focusing questions are on multi-level scales. We binarize them by assigning 1 as a label to examples where the majority of raters selected the top category (“high uptake”, “focusing question”) and 0 to all others.

4.2 Supervised Classification

We finetune RoBERTa (Liu et al., 2019) on turn-level annotations for each discourse feature. We run finetuning for 5 epochs, a batch size of 8 x 2 gradient accumulation steps. The choice of this model and parameters were optimal for efficient iteration on a single TitanX GPU; we leave extensive model exploration and hyperparameter search for future work.

For Student on Task and Student Reasoning, the input to the model was a single student utterance. For Teacher on Task, the input to the model was a single teacher utterance. For High Uptake and Focusing Question, the model input was a student utterance and a subsequent teacher utterance, to match what annotators saw while labeling. We balance labels during training by oversampling the minority category to represent 50% of labels, as we found this process yields better results.

4.3 Regression Analysis

Since our ultimate goal is to improve instruction via classroom discourse analysis, we need to understand if our NLP measures of discourse features

indeed correlate with observation scores and student outcomes. To understand this question, we (i) follow the procedure described above to fine-tune classifiers for each discourse on *all* annotations, (ii) predict discourse features for the entire NCTE transcript dataset, (iii) run regressions using classroom observation scores and value added scores as dependent variables.

Model. We run a linear regression, clustering standard errors at the teacher level. The models are captured by this equation:

$$y_d = x_f \beta_1 + T \beta_2 + S \beta_3 + \varepsilon \quad (1)$$

, where y_d is a vector representing a particular dependent variable d , x_f is a vector representing our predictions for a particular discourse feature f , T is a matrix of teacher covariates, S is a matrix of student covariates, $\beta_1, \beta_2, \beta_3$ are vectors of unknown parameters to be estimated and ε is a vector of residuals.

As for discourse features f , we use discourse moves in Table 3 and also include baseline measures of student talk ratios (% of student words and % of student turns) for comparison. We estimate the effect of these features on six different dependent variables d , which include five items from observation scores and value added scores.

Observation scores as dependent variables. To measure mathematics instructional quality, we use the main holistic item from the MQI instrument, a 5-level rating for lesson quality. The other four variables come from the CLASS scoring instrument: instructional dialogue, teacher sensitivity, teachers' regard for student perspectives and positive classroom climate. We picked these items *a priori*—before conducting the regressions—based on their relevance to dialogic instruction. One could choose other items from these observation protocols to conduct similar analyses.

Since observation scores are linked to transcripts, we first aggregate discourse move predictions to the lesson level. Specifically, we sum discourse feature predictions in each transcript, and divide them by the class duration (number of 7.5 minute segments in the transcript⁷).

Value added scores as dependent variables. Each teacher is linked to one value added score

⁷We do not have consistent timestamps that are more granular than 7.5 minutes.

per year given that these are based on end-of-the-year standardized test scores. Therefore, we mean-aggregate lesson-level data obtained above to the teacher-year level when conducting regressions with value-added scores.

Covariates. We include several covariates for teacher and classroom demographic. We include binary indicators for teacher gender and race/ethnicity and a numerical variable for years of experience. We include variables related to classroom composition in terms of student gender, race, free or reduced lunch status, special education status and limited English proficiency status. Table 1 includes a list of variables representing each of these categories.

5 Results

Supervised classification. Table 4 shows the performance of our supervised classifiers, averaged across five-fold cross validation. The results show that we can train our model to automatically classify each discourse feature with moderate to high accuracy. The model performs best on classifying on vs off task instruction — F1 score (harmonic mean of precision and recall) is .942 and .923 for student and teacher utterances, respectively. The model performs moderately well on higher-inference discourse moves, including Student Reasoning (F1 = .651), High Uptake (F1 = .688), and Focusing Questions (F1 = .501). We expect that model choice and hyperparameter tuning can improve the performance by 10-20%. However, even the human raters are only able to reach moderate agreement on these measures (Demszky et al., 2021b; Alic et al., 2022) and other analogous ones in classroom observation instruments (Kelly et al., 2020). The moderate human agreement indicates that these measures are subjective, which may set an upper bound to the models' performance.

Correlation with outcomes. Table 5 shows the correlation of each discourse feature with outcomes. We also include two baseline discourse features that measure student talk ratios: the percentage of student turns and the percentage of student words in each transcript. We find that **all discourse features** predicted by our classifiers correlate significantly with each classroom observation item measuring instruction quality and classroom climate.

All dialogic discourse moves — High Uptake, Focusing Questions and Student Reasoning — also

Measure	Accuracy	Precision	Recall	F1
Student on Task	0.902	0.952	0.931	0.942
Teacher on Task	0.867	0.932	0.914	0.923
Teacher Uptake	0.768	0.719	0.674	0.688
Focusing Question	0.856	0.474	0.538	0.501
Student Reasoning	0.863	0.644	0.666	0.651

Table 4: Performance of our RoBERTa classifier on each discourse feature. Values are averages across a five-fold cross validation.

	(1)	(2)	(3)	(4)	(5)	(6)
	Value Added Scores	Math Instruction Quality	Instructional Dialogue	Teacher Sensitivity	Regard for Student Perspectives	Positive Climate
Student on Task	0.038+ (0.020)	0.022* (0.010)	0.032** (0.011)	0.033** (0.008)	0.024** (0.007)	0.036** (0.007)
Teacher on Task	0.038+ (0.020)	0.021* (0.010)	0.030** (0.010)	0.034** (0.008)	0.024** (0.007)	0.035** (0.007)
Teacher Uptake	0.121* (0.050)	0.117** (0.032)	0.083** (0.026)	0.089** (0.019)	0.058** (0.017)	0.079** (0.017)
Focusing Question	0.234* (0.104)	0.233** (0.086)	0.198** (0.072)	0.132** (0.035)	0.164** (0.044)	0.115** (0.036)
Student Reasoning	0.191* (0.091)	0.313** (0.066)	0.246** (0.050)	0.144** (0.031)	0.173** (0.035)	0.120** (0.035)
Student Turn %	1.044 (1.357)	-0.047 (0.528)	0.718 (0.669)	0.214 (0.574)	0.125 (0.485)	-0.172 (0.560)
Student Word %	0.359 (0.792)	0.721+ (0.413)	1.132* (0.541)	0.001 (0.325)	0.469 (0.395)	0.322 (0.387)
Observations	523	1557	1554	1554	1554	1554

Table 5: The correlation of discourse features with outcomes, estimated at the transcript level. Each cell represents coefficients from a separate regression. Standard errors are enclosed in parentheses. Dependent variables are standardized. The ** $p < 0.01$, * $p < 0.05$, + $p < 0.1$

show a significant positive correlation with teachers’ value added scores. Specifically, each additional High Uptake per 7.5 minute segment increases teachers’ value added scores by 12% of a standard deviation. Analogically, each additional instance of a Focusing Question and Student Reasoning per 7.5 minute segment increases teachers’ value added scores by 23% and 19% of a standard deviation, respectively. Student on Task and Teacher on Task also correlate with marginal significance with value added scores.

Interestingly, baseline measures of student talk percentages do not in themselves correlate with observation scores or value added scores. One exception is student word percentage, which correlates positively with ratings for instructional dialogue.

6 Discussion

We find that our measures of discourse moves, which we identified by consulting research on mathematics instruction, correlate with human raters’ perceptions of lesson quality, and with students’ learning outcomes. These results are significant in multiple ways.

6.1 Implications & Significance

The fact that all of our measures correlate with MQI and CLASS observation scores indicate that the automated measures align with expert evaluations of instruction. This is a key result that provides external validation for these automated measures. The positive correlation of teacher value-added scores with measures for teacher uptake, focusing questions and student reasoning suggests that the use of these dialogic talk moves is associated with student

learning. Substantively, this finding contributes evidence that classrooms where students are more deeply engaged with mathematical ideas — and where teachers use their students’ mathematical contributions — are more likely to produce better achievement outcomes (O’Connor and Michaels, 1993; Michaels and O’Connor, 2015).

These findings become even more significant in the context of baseline measures of student engagement — percentage of student turns and percentage of student words — which do not show positive correlations with instruction quality and value added scores. These findings add to a collection of mixed results by related work, some of which show positive, some of which show no relationship between student talk time and learning outcomes (see Sedova et al., 2019, for an overview).

6.2 Limitations

Unmeasured covariates. There is a wide range of factors that affect instructional outcomes, only a subset of which could be measured with this data. Making strong claims about the link between discourse moves and instructional outcomes requires experimental validation. For example, the quality of the math task that the students are working may affect the discourse as well as learning outcomes. We can isolate the effect of discourse moves by randomly assigning teachers to learning opportunities that successfully help them improve their use of these moves, and examining downstream impacts of these new talk moves on student outcomes. (Demszky et al., 2021a) has taken a similar approach successfully in an informal teaching context, but such a study is yet to be done in a K-12 context.

Generalizability. Although the NCTE transcript dataset is the largest available dataset of U.S. classroom transcripts, it only captures a tiny fraction of U.S. classrooms and hence there are limitations to its representativeness. The data represents mostly white female teachers working in mid-size to large districts, so it would be valuable to collect new data from other types of districts and a more diverse teacher population. The fact that the data was collected a decade ago may pose limitations to its ongoing relevance; during the period under study (2010-2013), many schools were transitioning toward Common Core-aligned instruction in mathematics but yet lacked high-quality curriculum materials for doing so. That said, research in

education reform has long attested to the fact that teaching practices have remained relatively constant over time (Cuban, 1993; Cohen and Mehta, 2017) and that there are strong socio-cultural pressures that maintain the instructional status quo (Cohen, 1988). In general, it is important to carefully validate measures built on the NCTE data on a new domain to ensure that it is representative of the target population.

Limitations of linked data. Education research has attested the limitations of standardized assessments in capturing student learning and reasoning (Sussman and Wilson, 2019). Student questionnaires in the NCTE data can provide an alternative perspective on students’ experiences and mathematics outcomes but these responses have a lot of missing values, and hence it may not provide robust estimates. Furthermore, understanding equity in instruction is a high priority for our research team and for the field more generally. However, studying equity within this data is challenging, since student speakers are not linked to administrative files containing student background and achievement variables. That said, such speaker-level demographic data is rarely available in instructional contexts, for important ethical reasons, and thus this limitation may encourage researchers to develop measures of instructional equity that leverage classroom-level, instead of speaker-level demographic information.

6.3 Ethical Considerations

We outline measures to safeguard students and teachers in this data and to ensure that the data and those potentially impacted by downstream uses of this data.

Consent & privacy. Both parents and teachers gave consent for the de-identified data to be retained and used in future research (Harvard IRB #17768). It is our highest priority that the identity of teachers and students in the data are kept private. Given that the none of the district and school names are disclosed, and that transcripts are fully de-identified, it is not possible to recover the identity of teachers and students in this data.

Representation. As Madaio et al. (2022) point out, applying AI in the educational domains comes with a risk of propagating and exaggerating existing inequities. As we describe in the paper, the data represents a largely low-income, demographically diverse student population. This means that

the data can help with creating measures that are representative of low income students, students of color and students who are English language learners who receive the type of instruction captured in this dataset. As we point out above, the data should not be assumed to represent the diversity of identities and experiences of all students and teachers in U.S. classrooms and different forms of instruction. Furthermore, the data was annotated by raters whose demographics are largely representative of teacher demographics in the US⁸ (Demszky et al., 2021b), which, just like in this data, does not unfortunately match U.S. student demographics. Rater bias (Campbell and Ronfeldt, 2018) cannot be outruled especially given the subjectivity of these constructs.

Downstream application. Users of this data have to agree to never use the data in a way that may cause harm to students and teachers, such as to build tools that discriminate against different groups of students and teachers or to surveil and punish teachers based on their practice. The sole purpose of this data should be to help us understand and facilitate student-centered and equitable instructional practice, and to empower historically marginalized teacher and student populations.

6.4 Directions for Future Work

The NCTE dataset opens up numerous directions for future work, some of which we are currently pursuing. First, one key direction is **building new NLP measures** of instruction. Observation protocols such as MQI, CLASS and the Culturally Responsive Instruction Observation Protocol (CRIOP) (Powell et al., 2016) and related work by Suresh et al. (2022) and Hunkins et al. (2022) can provide inspiration for various discourse features that can be measured in this data. Given the context-dependence of several instructional moves, new measures can incorporate more context beyond a single utterance or exchange between the student and the teacher. It would also be valuable to incorporate lesson-level metadata in the NLP model, e.g. lesson keywords, date, grade level, to create more context-specific measures.

One can also conduct **bottom-up exploration of linguistic patterns** in the data to inform education research. For example, one could create an equity gap measure, leveraging academic outcomes and

⁸<https://nces.ed.gov/fastfacts/display.asp?id=28>

classroom demographics, and compare transcripts from classrooms with low equity gap with ones from classrooms with high equity gap. Doing so can help identify instructional correlates of equity gap, and help us understand how we can facilitate equitable instructional practices.

Third, since no two education settings are the same, it would be extremely valuable to **complete the NCTE dataset with other educational datasets** from a diverse range of settings, including various school subjects, informal and formal, online and in person, group and one-on-one settings, as well as data from multiple regions and countries and from multiple modalities, including text, audio and video. These datasets together can help us understand how effective instruction looks like across teaching contexts and modalities and build measures sensitive to these contextual differences.

Finally, one can **apply what we learn from this data to improve instruction**. For example, measures built on this data can power automated feedback tools (Suresh et al., 2021; Demszky et al., 2021a) and enable empirically-driven improvements to widely used professional development frameworks (e.g. Gregory et al., 2017, NCTM guides⁹). We hope that this resource can power efforts to address pressing issues in education, such as pandemic learning loss and equity gaps in mathematics instruction.

7 Conclusion

We introduce a dataset of elementary math classroom transcripts, associated with a rich source of linked variables. We train classifiers on turn-level annotations to predict discourse moves and show that predictions for these moves correlate significantly with observation scores and value added scores. These results demonstrate how the NCTE dataset can serve as a valuable resource for understanding classroom interactions and for powering tools that seek to facilitate instruction.

Acknowledgments

We thank Jing Liu, Hannah Kleen, Shyamoli Sanghi, Zach Himmelsbach, Zid Mancenido and Dan Jurafsky for the productive conversations and their collaboration. We are grateful to Dan McGinn for help with data sharing and to Christine Kuzdzal and Carol DeFreese for help with annotating student reasoning.

⁹<https://www.nctm.org/pdguides/>

References

- Robin John Alexander. 2008. *Towards Dialogic Teaching: rethinking classroom talk (4th Edition)*. Dialogos.
- Sterling Alic, Dorottya Demszky, Zid Mancenido, Jing Liu, Heather Hill, and Dan Jurafsky. 2022. Computationally identifying funneling and focusing questions in classroom discourse. In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 224–233.
- M. M. Bakhtin. 1981. *The dialogic imagination: four essays*. University of Texas Press.
- McKenzie Braley and Gabriel Murray. 2018. The group affect and performance (gap) corpus. In *Proceedings of the ICMI 2018 Workshop on Group Interaction Frontiers in Technology (GIFT)*.
- Jere E Brophy. 1984. *Teacher behavior and student achievement*. 73. Institute for Research on Teaching, Michigan State University.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Shanyce L Campbell and Matthew Ronfeldt. 2018. Observational evaluation of teachers: Measuring more than we bargained for? *American Educational Research Journal*, 55(6):1233–1267.
- David K Cohen. 1988. *Teaching practice: Plus ça change*. National Center for Research on Teacher Education East Lansing, MI.
- David K Cohen and Jal D Mehta. 2017. Why reform sometimes succeeds: Understanding the conditions that produce reforms that last. *American Educational Research Journal*, 54(4):644–690.
- Larry Cuban. 1993. *How teachers taught: Constancy and change in American classrooms, 1890-1990*. Teachers College Press.
- Dorottya Demszky, Jing Liu, Heather C Hill, Dan Jurafsky, and Chris Piech. 2021a. Can automated feedback improve teachers’ uptake of student ideas? evidence from a randomized controlled trial in a large-scale online course. edworkingpaper no. 21-483. *Annenberg Institute for School Reform at Brown University*.
- Dorottya Demszky, Jing Liu, Zid Mancenido, Julie Cohen, Heather Hill, Dan Jurafsky, and Tatsunori B Hashimoto. 2021b. Measuring conversational uptake: A case study on student-teacher interactions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1638–1653.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- John J Godfrey and Edward Holliman. 1997. Switchboard-1 release 2. *Linguistic Data Consortium, Philadelphia*, 926:927.
- A Gregory, E Ruzek, CA Hafen, A Yee Mikami, JP Allen, and RC Pianta. 2017. My teaching partner-secondary: A video-based coaching model. *Theory into practice*, 56(1):38–45.
- Heather C Hill, Merrie L Blunk, Charalambos Y Charalambous, Jennifer M Lewis, Geoffrey C Phelps, Laurie Sleep, and Deborah Loewenberg Ball. 2008. Mathematical knowledge for teaching and the mathematical quality of instruction: An exploratory study. *Cognition and instruction*, 26(4):430–511.
- Nicholas Hunkins, Sean Kelly, and Sidney D’Mello. 2022. “beautiful work, you’re rock stars!”: Teacher analytics to uncover discourse that supports or undermines student motivation, identity, and belonging in classrooms. In *LAK22: 12th International Learning Analytics and Knowledge Conference*, pages 230–238.
- Daniel Jurafsky, Elizabeth Shriberg, and Debra Bisca. 1997. Switchboard SWBD-DAMSL Labeling Project Coder’s Manual, Draft 13. Technical Report 97-02, University of Colorado Institute of Cognitive Science.
- T Kane, H Hill, and D Staiger. 2015. National center for teacher effectiveness main study. icpsr36095-v2.
- Thomas J Kane, Daniel F McCaffrey, Trey Miller, and Douglas O Staiger. 2013. Have we identified effective teachers? validating measures of effective teaching using random assignment. research paper. met project. *Bill & Melinda Gates Foundation*.
- Sean Kelly, Robert Bringe, Esteban Aucejo, and Jane Cooley Fruehwirth. 2020. Using global observation protocols to inform research on teaching effectiveness and school improvement: Strengths and emerging limitations. *Education Policy Analysis Archives*, 28:62.
- Sean Kelly, Andrew M Olney, Patrick Donnelly, Martin Nystrand, and Sidney K D’Mello. 2018. Automatically measuring question authenticity in real-world classrooms. *Educational Researcher*, 47(7):451–464.

- Jing Liu and Julie Cohen. 2020. Measuring teaching practices at scale: A novel application of text-as-data methods. *EdWorking PaperNo*, pages 20–239.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Michael Madaio, Su Lin Blodgett, Elijah Mayfield, and Ezekiel Dixon-Román. 2022. Beyond “fairness”: Structural (in) justice lenses on ai for education. In *The Ethics of Artificial Intelligence in Education*, pages 203–239. Routledge.
- Sarah Michaels and Catherine O’Connor. 2015. Conceptualizing talk moves as tools: Professional development approaches for academically productive discussion. *Socializing intelligence through talk and dialogue*, 347:362.
- Sarah Michaels, Mary Catherine O’Connor, Megan Williams Hall, and Lauren B Resnick. 2010. Accountable talk sourcebook: For classroom conversation that works. *Pittsburgh, PA: University of Pittsburgh Institute for Learning*.
- Martin Nystrand, Adam Gamoran, Robert Kachur, and Catherine Prendergast. 1997. *Opening dialogue*. New York: Teachers College Press.
- Mary C O’Connor and Sarah Michaels. 1993. Aligning academic task and participation status through revoicing: Analysis of a classroom discourse strategy. *Anthropology & Education Quarterly*, 24(4):318–335.
- James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates*, 71(2001):2001.
- Robert C Pianta, Karen M La Paro, and Bridget K Hamre. 2008. *Classroom Assessment Scoring System: Manual K-3*. Paul H Brookes Publishing.
- Rebecca Powell, Susan Chambers Cantrell, Victor Malo-Juvera, and Pamela Correll. 2016. Operationalizing culturally responsive instruction: Preliminary findings of criop research. *Teachers College Record*, 118(1):1–46.
- Lynn C Reimer, Katerina Schenke, Tutrang Nguyen, Diane K O’dowd, Thurston Domina, and Mark Warschauer. 2016. Evaluating promising practices in undergraduate stem lecture courses. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2(1):212–233.
- Lorenzo A. Rossi and Omprakash Gnawali. 2014. Language Independent Analysis and Classification of Discussion Threads in Coursera MOOC Forums. In *Proceedings of the IEEE International Conference on Information Reuse and Integration (IRI 2014)*.
- Katherine Schaeffer. 2021. [America’s public school teachers are far less racially and ethnically diverse than their students.](#)
- Daniel L Schwartz, Catherine C Chase, Marily A Oppezzo, and Doris B Chin. 2011. Practicing versus inventing with contrasting cases: The effects of telling first on learning and transfer. *Journal of educational psychology*, 103(4):759.
- Klara Sedova, Martin Sedlacek, Roman Svaricek, Martin Majcik, Jana Navratilova, Anna Drexlerova, Jakub Kychler, and Zuzana Salamounova. 2019. Do those who talk more learn more? the relationship between student classroom talk and student achievement. *Learning and instruction*, 63:101217.
- Pooja G Sidney, Shanta Hattikudur, and Martha W Alibali. 2015. How do contrasting cases and self-explanation promote learning? evidence from fraction division. *Learning and Instruction*, 40:29–38.
- Katherine Stasaski, Kimberly Kao, and Marti A. Hearst. 2020. [CIMA: A large open access dialogue dataset for tutoring.](#) In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–64, Seattle, WA, USA → Online. Association for Computational Linguistics.
- Abhijit Suresh, Jennifer Jacobs, Charis Clevenger, Vivian Lai, Chenhao Tan, James H Martin, and Tamara Sumner. 2021. Using ai to promote equitable classroom discussions: The talkmoves application. In *International Conference on Artificial Intelligence in Education*, pages 344–348. Springer.
- Abhijit Suresh, Jennifer Jacobs, Charis Harty, Margaret Perkoff, James H Martin, and Tamara Sumner. 2022. The talkmoves dataset: K-12 mathematics lesson transcripts annotated for teacher and student discursive moves. *arXiv preprint arXiv:2204.09652*.
- Abhijit Suresh, Tamara Sumner, Jennifer Jacobs, Bill Foland, and Wayne Ward. 2019. Automating analysis and feedback to improve mathematics teachers’ classroom discourse. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 9721–9728.
- Joshua Sussman and Mark R Wilson. 2019. The use and validity of standardized achievement tests for evaluating new curricular interventions in mathematics and science. *American Journal of Evaluation*, 40(2):190–213.
- Gordon Wells. 1999. *Dialogic inquiry: Towards a socio-cultural practice and theory of education*. Cambridge University Press.