

# Measuring Teaching with LLMs

Michael Hardy  
Stanford University  
hardym@stanford.edu

## Abstract

Objective and scalable measurement of teaching quality is a persistent challenge in education. While Large Language Models (LLMs) offer potential, general-purpose models have struggled to reliably apply complex, authentic classroom observation instruments. This paper uses custom LLMs built on sentence-level embeddings, an architecture better suited for the long-form, interpretive nature of classroom transcripts than conventional subword tokenization. We systematically evaluate five different sentence embeddings under a data-efficient training regime designed to prevent overfitting. Our results demonstrate that these specialized models can achieve human-level and even super-human performance with expert human ratings above 0.65 and surpassing the average human-human rater correlation. Further, through analysis of annotation context windows, we find that more advanced models—those better aligned with human judgments—attribute a larger share of score variation to lesson-level features rather than isolated utterances, challenging the sufficiency of single-turn annotation paradigms. Finally, to assess external validity, we find that aggregate model scores align with teacher value-added measures, indicating they are capturing features relevant to student learning. However, this trend does not hold at the individual item level, suggesting that while the models learn useful signals, they have not yet achieved full generalization. This work establishes a viable and powerful new methodology for AI-driven instructional measurement, offering a path toward providing scalable, reliable, and valid feedback for educator development.<sup>1</sup>

## 1 Introduction

Measuring teaching quality is hard (Jurenka et al., 2024; Kane and Staiger, 2012; Ho and Kane,

<sup>1</sup>[https://github.com/hardy-education/measuring\\_teaching\\_encoders](https://github.com/hardy-education/measuring_teaching_encoders)

2013). Despite their ubiquity as the primary form of teacher development and evaluation, ratings of instructional quality, even when conducted by trained experts, have low reliability, unknown accuracy, and are very expensive to conduct (Ho and Kane, 2013; Kane et al., 2015; Kane and Staiger, 2012; Glaese et al., 2022; Whitehill and LoCasale-Crouch, 2024; Whitehurst et al., 2014; Jurenka et al., 2024; Tack et al., 2023; Grissom et al., 2013; Liu and Cohen, 2021). Recent work has sought to reduce the costs of these evaluations using large language models (LLMs) to annotate spoken discourse in classrooms to support such evaluations on actual instruments used with educators, but such models have not yet shown the ability to help with these complex tasks (Wang and Demszky, 2023; Whitehill and LoCasale-Crouch, 2024; Xu et al., 2024; Hardy, 2024). This study builds on these studies, answering calls to do more to evaluate the capacity of LLMs for classroom tasks (Casabianca et al., 2013; Liu and Cohen, 2021).

This study investigates whether pretrained contextual embeddings at the sentence level can meaningfully capture classroom dialogue for automated assessment of instructional quality. We systematically evaluate multitask encoder models trained on fixed sentence-level embeddings to predict expert human ratings of teaching effectiveness across 25 distinct instructional dimensions. We achieve state-of-the-art performance on this task, surpassing existing human benchmark correlations while providing novel insights into the training dynamics of multitask models applied to educational discourse.

Our findings have significant implications for educational assessment and the broader application of NLP methods to specialized domains. We show that shared-weight multitask architectures initially learn general representations of teaching quality that align well with student outcomes, but continued training may lead to overfitting to noisy human

annotations on individual constructs rather than the underlying pedagogical constructs of interest. These insights suggest new directions for developing robust AI systems for classroom analysis and highlight fundamental challenges in aligning automated assessments with meaningful educational outcomes.

## 1.1 Primary Contributions

1. **A longitudinal analysis of representation learning for teaching quality.** We provide the first systematic evidence of how different sentence-embedding LLMs develop an understanding of effective instruction throughout their training process.
2. **A new benchmark for automated instructional rating.** We demonstrate state-of-the-art performance, outperforming existing models by achieving the highest human-expert correlation across 25 distinct instructional dimensions.
3. **A critique of single-turn evaluation.** Through the first analysis of score stability over time, we show that static, single-turn evaluations are insufficient and introduce a more robust, temporally-aware method for assessing LLM-based ratings.
4. **A validation framework linking ratings to student achievement.** We introduce the first methodology to directly measure the alignment between an LLM’s ratings of teaching and externally-validated measures of student learning gains (teacher value-added).

## 2 Related Work and Motivation

### 2.1 Sentence Embeddings

The robustness of pre-trained language models to out-of-distribution text remains an open question, particularly for specialized domains such as educational settings where child speech patterns and pedagogical discourse structures predominate. To address this challenge, we investigate whether sentence-level embeddings (Reimers and Gurevych, 2019) can provide more stable representations of classroom language than traditional subword tokenization approaches, despite the potential loss of fine-grained semantic information. We test the large versions of each of the following pre-trained sentence embeddings: Unsupervised SimCSE and Supervised SimCSE (Gao et al.,

2022), E5 (Wang et al., 2024a), Multilingual E5 (Wang et al., 2024b), GTE (Li et al., 2023), and a contrast fine-tuned RoBERTa model released with sentence-transformers (Liu et al., 2019; Reimers and Gurevych, 2019). This diversity in embedding approaches also allows us to investigate potential biases in model interpretation across different linguistic communities and teaching contexts.

### 2.2 Teacher Development and Evaluation

School leaders working with teachers to improve the quality of instruction typically evaluate the teacher’s proficiency in a range of competencies (typically measured during in-class observation and evaluation on a teaching rubric; (Aguilar, 2013; Bambrick-Santoyo, 2016, 2018)), a process that is often time-consuming and produces ratings (labels) that are unreliable (Kane and Staiger, 2012; Blazar, 2018; Kane et al., 2013; Casabianca et al., 2013). Without accurate classifications, it is challenging for practitioners to prioritize instructional needs and aligned practices from among the many elements of good teaching (Saphier et al., 2008; Darling-Hammond, 2014; Hammond, 2015; Lemov and Atkins, 2015; Lemov, 2021; Liljedahl et al., 2021; Darling-Hammond et al., 2020; Schwartz et al., 2016) and for researchers to empirically quantify the impact of good teaching practices (Pianta and Hamre, 2009; Charalambous and Delaney, 2019; Blazar and Polard, 2022; Jurenka et al., 2024). Thus, this work provides a bridge to research seeking to improve teaching quality by providing feedback to teachers on various instructional techniques (Samei et al., 2014; Donnelly et al., 2017; Kelly et al., 2018; Demszky et al., 2021; Suresh et al., 2022; Jacobs et al., 2022; Alic et al., 2022; Demszky and Liu, 2023; Demszky et al., 2024, 2023).

**Automated Evaluation** Several studies have investigated automated evaluation (Whitehill and LoCasale-Crouch, 2024; Wang and Demszky, 2023; Xu et al., 2024; Hardy, 2024). This study builds on these studies, and replicates the encoder model constructions described by Hardy.<sup>2</sup>

## 3 Methods

For each method, we also display the results at the item-level for better understanding of the learning process.

<sup>2</sup>A more complete description of the model architecture is in (Hardy, 2024)

### 3.1 Human Expert Rating Correlation

To find a generalized measure of correlation across items similar to CLASS and MQI, we use a multi-level partial Spearman’s correlation with inference based on item-level random effects. This accounts for the hierarchical structure while providing a robust, rank-based measure of association that generalizes beyond the specific items sampled.

$$\rho_{\text{part}} = \text{Corr}(\tilde{R}_{ij}^{(1)}, \tilde{R}_{ij}^{(2)} \mid \mathbf{J}_{ij})$$

### 3.2 LLM Rating Stability via Variance Decomposition

We employed a generalizability theory framework (Brennan, 2001) to decompose variance in automated LLM scoring across six nested hierarchical levels: sentences (X) within utterances (U) within chapters (C) within lesson stages (S) within lessons (L) within teachers (T), denoted as  $X : U : C : S : L : T$ . This design enables quantification of context dependency as models evolve during training. The proportion of variance attributable to each level,  $h \in \{T, L, S, C, U, X, e\}$ , is:

$$\rho_h = \frac{\sigma_h^2}{\sigma_T^2 + \sigma_L^2 + \sigma_S^2 + \sigma_C^2 + \sigma_U^2 + \sigma_X^2 + \sigma_e^2}$$

### 3.3 External Validity via $\tau$ -Canonical Correlation Analysis

We employ canonical correlation (Hotelling, 1936) analysis (CCA) with a Kendall’s tau (Kendall, 1938, 1945) kernel to measure alignment between teacher value-added measures (VAMs) and classroom instructional ratings. In this case, we need a metric that captures the directional alignment only, as differences in scales and ranks may not be meaningful.<sup>3</sup> We briefly reconstruct the  $\tau$  kernel for creation of scatter matrices here to motivate it as a highly robust (Bishara and Hittner, 2017) measure of alignment between LLM ratings,  $\mathbf{x}$ , and another metric,  $\mathbf{y}$  and as having a straightforward alignment interpretation. We translate the following statement into the desired kernel: *"LLM X rates lesson A as better than lesson B:  $[x_A > x_B]$ . Does the order align with student learning results Y associated*

<sup>3</sup>For example, if two raters give Lesson A the same score of 7, but give Lesson B different ratings, 3 and 4, we would not have evidence to support the notion that LLMs (or even humans) apply an instructional rubric precisely enough for such differences in interval ranks to be practically meaningful when measuring alignment.

with each lesson,  $[y_A > y_B]$ ?" Thus, for any two lessons, indexed by  $i$  and  $j$  and with brackets as indicator functions, we construct this relationship for each:

$$x_{ij} = [j > i] - [j < i], \quad y_{ij} = [j > i] - [j < i]$$

We construct Gram matrices  $\mathbf{K}_X$  and  $\mathbf{K}_Y$  where  $[\mathbf{K}_X]_{ij} = K_\tau(\mathbf{x}_i, \mathbf{x}_j)$  and analogously for  $K_Y$ , which are used to solve the eigenvalue problem required for canonical correlation. In our case, all matrices were positive semi-definite and no smoothing was needed. The results are robust to typical transformations for the calculation.<sup>4</sup>

## 4 Data and Experiment

### 4.1 Data

The original classroom lessons used in this study are from the National Center for Teacher Effectiveness (NCTE) Main Study (Kane et al., 2015),<sup>5</sup> which contains 3 years of data collection and observations of math instruction in approximately 50 schools and three hundred (4th and 5th grade) mathematics classrooms across 4 school districts in the United States. This rich dataset contains **multiple measures of teaching effectiveness**, including expert ratings of the lessons classrooms, and **multiple high-quality measures student learning gains**.

**Classroom Lessons and Text** Human raters watched videos classrooms, and the transcripts<sup>6</sup> of these same videos (Demszky and Hill, 2022) are used by LLMs for the same task, where the class discourse is equipartitioned across words to align the text with human labels in the absence of timestamps, following (Hardy, 2024).

**Observation Instruments** Our approach encompasses two complementary observation frameworks: a 12-item general teaching practices instrument and a 13-item mathematics-specific teaching assessment (Bacher-Hicks et al., 2017, 2019): the CLASS framework (Pianta et al., 2008) for general instructional practice and the content-specific Mathematical Quality of Instruction (MQI) (Hill

<sup>4</sup>We also compute the generalized eigenvalue problem using the methods put forward by (Yoon et al., 2020) with the results in the appendix.

<sup>5</sup><https://www.icpsr.umich.edu/web/ICPSR/studies/36095/datadocumentation>

<sup>6</sup><https://github.com/ddemszky/classroom-transcript-analysis>

et al., 2008). The 63 MQI raters and the 19 external CLASS raters attended biweekly calibration meetings to ensure standardization of scoring procedures. Both frameworks are composed of multiple items that represent distinct instructional dimensions to be evaluated (Hill et al., 2008; Hardy, 2024; Hill et al., 2012; Kane and Staiger, 2012). The MQI and CLASS also represent two types of task for LLMs—detection and summarization, respectively—a distinction that is also clearly illustrated by the distributions of scores coming from the human raters (see Figure 2). Human rating experts watched videos and provided ratings on all MQI and CLASS items at regular intervals throughout the class, resulting in 779,107 unique numeric ratings provided in 1,762 lessons<sup>7</sup> delivered by 317 teachers to more than 10,000 students in 53 schools.

**Value-added measures (VAMs)** are the current gold standard for estimating teacher effects on student achievement gains (Kane and Staiger, 2012; Bacher-Hicks et al., 2017, 2019). VAMs use prior student achievement data and other covariates to measure whether a student’s end-of-year performance was above or below the student’s expected performance on a standardized exam. The teacher-level VAM is an estimate of the combined deviations from expected performance for their students, offering a rigorous aggregated estimate of a teacher’s contribution to student learning gains over a school year. As far as we are aware, this is the first study to test LLMs using standardized and value-added measures of student learning. Rare for education datasets, the data of the present study have multiple VAM measurements, which we will use together as a random vector for canonical correlation (and stacked at the year level (Bacher-Hicks et al., 2019) for item-level comparisons).

## 4.2 Encoder Model Construction

We develop custom encoder architectures based on sentence-level embeddings to address four key research questions: (1) embedding efficiency in model training, (2) performance relative to human raters, (3) score variation across different temporal contexts, and (4) alignment with rigorous measures of student learning outcomes.

**Why Sentence Embeddings? A More Interpretable Architecture.** Analyzing lengthy class-

<sup>7</sup>Transcripts are available for 1,600 of the lessons (Demszky and Hill, 2022)

room discourse poses a challenge for standard LLMs, whose subword tokens are often too granular and computationally intensive for such long-form text. To overcome this, we chose sentence-level embeddings as our foundation. Sentences are natural, interpretable semantic units, allowing our models to efficiently process entire lessons and directly map what a teacher says to established pedagogical frameworks. We evaluated five pretrained sentence-embedding models within a multi-task architecture that learns to score 25 distinct teaching dimensions simultaneously. The model’s core weights are shared across all tasks, reflecting the pedagogical theory that effective teaching is a cohesive skill set. Only the final output layers are specialized for each dimension, allowing the model to learn both general and specific features of instructional quality.

**Structuring Data for Meaningful Analysis.** To capture the natural rhythm of a class, we organized transcripts into a three-phase structure (beginning, middle, end) common in elementary math. Each phase was then divided into chapters of a fixed duration (7.5 minutes for MQI, 15 for CLASS). We processed teacher speech into individual sentences, creating uniform inputs and enabling the model to update its assessment with each new utterance. We further augmented our dataset using a sliding-window technique to generate additional training samples from the discourse.

**An Efficient and Robust Training Protocol.** Given the challenge of collecting high-quality observation data, our training protocol prioritized efficiency and generalization. All models were trained for five epochs. We first determined the optimal optimizer by testing AdamW and Adamax; Adamax proved more stable for four of the five models under the strong regularization needed to prevent overfitting. The fifth model (GTE) performed better with AdamW. We retained all models to ensure a comprehensive evaluation. Additional training details follow the protocol in (Hardy, 2024).

## 5 Discussion

### 5.1 Some embeddings are better than others

Not all contextual embeddings capture the same semantic information. The RoBERTa, Unsupervised SimCSE, and E5 Multilingual models consistently outperformed other embedding models in the present study. When correlations are measured by




Figure 1: Correlation with human experts across training epochs. The MQI instrument had at least two human raters per lesson, and the mean and interquartile range of all 63 human MQI raters correlated across the other raters are represented by the gray line and shaded region in the figure.

aggregating to the lesson-level, all models’ ratings would be in the top quartile of human raters. The item level performance at the chapter level can be investigated more deeply in Fig 6.

**Implications** Future work should explore fine-tuning SimCSE embeddings or similar self-supervised fine-tuning techniques on related classroom transcripts to investigate the extent to which providing domain-specific contrastive learning could further capture the most relevant semantic information from classroom discourse. Sentence-level embeddings also provide a pathway to interpretability with feature attribution via integrated gradients. Usefulness of feature attribution methods rely on how interpretable each input feature is. In the case of classroom instruction, a sentence spoken is a meaningful unit of discourse, whereas more common methods of creating features for LLMs typically rely on subword tokenization, producing a pixelated semanticity much harder for humans to interpret.

## 5.2 Mature models show score stability across longer time windows

Analysis across training epochs revealed a systematic shift in variance attribution: early-epoch models exhibited substantial utterance-level variation ( $\rho_U$ ), while mature models demonstrated increased

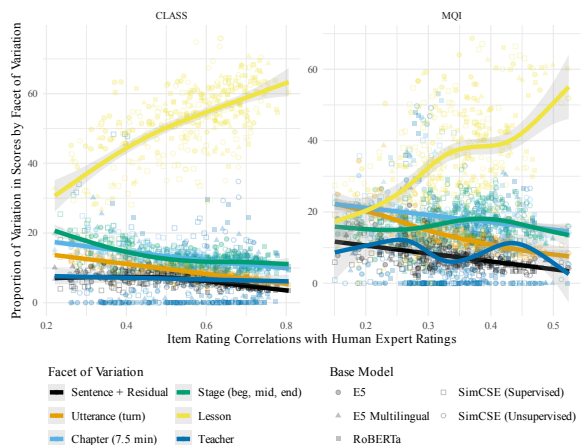


Figure 2: Proportion of variation explained as related to a model’s alignment to human expert ratings.

lesson-level variation ( $\rho_L$ ) with reduced local context dependency. This finding challenges prevailing single-turn annotation paradigms in LLM evaluation, suggesting that as models improve, they capture broader pedagogical patterns rather than utterance-specific features. Consequently, evaluation frameworks must incorporate extended conversational contexts and hierarchical sampling strategies to accurately assess model performance in educational applications.

We find that the variation in scores for models that are more aligned with human ratings, tends to be less sensitive to smaller changes in time. For the CLASS rubric in particular, more human-aligned models maintain more score stability across an entire lesson, suggesting that the scores are more representative of persistent differences in lessons.

## 5.3 External Validity

In order to assess for overfitting to this particular subset of human raters, we measured the alignment of the chapter-level multi-task scores against the value-added to learning for the students in the class. To date, there is no study of which we are aware that connect LLM measures of teaching to the actual value-added in student learning external validity.

We find that for is increasing alignment between VAM and classroom ratings as models mature and become more aligned with humans.

We demonstrate that while individual task performance generally improves with training, the correlation between model predictions and student achievement outcomes follows a non-monotonic trajectory. This phenomenon reveals important ten-




Figure 3:  $\tau$ -canonical correlation between classroom observation ratings and value-added measures as a function of model alignment to human expert ratings.

sions in multitask learning: as models become more specialized at distinguishing between specific instructional skills, their ability to capture the general teaching effectiveness that correlates with student learning gains may paradoxically diminish.

To assess whether our models exhibit overfitting to this particular subset of human raters, we examined the relationship between chapter-level multitask scores and value-added measures (VAM) of student learning outcomes. We employed  $\tau$ -canonical correlation analysis to quantify the strength of association between these two sets of variables while accounting for their multivariate nature.

This analysis addresses a critical gap in the literature: to our knowledge, no prior study has investigated the connection between LLM-derived measures of teaching quality and externally validated VAMs. Such validation is essential for establishing the practical utility of automated classroom assessment tools (Figure 3).

While the  $\tau$ -canonical correlation between model predictions and student achievement outcomes generally improves with model sophistication, individual item performance can follow an inverted trajectory performance (Figure 9). This phenomenon illuminates an important tension inherent in multitask learning for educational assessment. As models become increasingly specialized at distinguishing between specific instructional skills that human raters prioritize, their capacity to

capture the broader dimensions of teaching effectiveness that correlate with actual student learning gains may paradoxically diminish. This suggests that perfect alignment with human expert ratings may not constitute the optimal objective for developing classroom observation tools intended to predict student outcomes.

## 6 Conclusion

Rating classroom teaching quality remains a persistent challenge, with both human evaluators and large language models (LLMs) struggling to effectively utilize authentic classroom observation instruments. While general-purpose GPT models have shown limited promise for this task, we developed custom LLMs based on sentence embeddings to overcome the interpretability and scalability limitations of traditional subword tokenization approaches when processing lengthy classroom transcripts. We systematically evaluated five pretrained sentence embedding models using identical training regimes designed to maximize efficiency and minimize overfitting given the scarcity of authentic classroom data. We assessed their ability to capture pedagogically relevant information using established observation frameworks. Our results demonstrate that three embedding models achieved human-level performance, with correlations exceeding 0.65 for CLASS and surpassing human averages for MQI. More mature models increasingly attribute variation to differences at lesson-level rather than utterance-specific features. This finding challenges prevailing single-turn evaluation paradigms in LLM assessment and development, suggesting that improved models capture broader pedagogical patterns in long-context classroom dynamics. Validity analysis using teacher value-added measures revealed that while models achieving better human alignment also showed stronger alignment with learning outcomes in aggregate, this relationship did not hold at the item level. These results indicate that although models learn pedagogically meaningful features, evidence for generalization remains limited, highlighting important directions for future development of automated teaching quality assessment systems.

## Limitations

The findings of this study should be considered in light of several limitations related to the data, models, and readiness for practical application. We

position this work as a proof of concept, and the following factors must be addressed before these methods can be considered for real-world implementation.

**Scope of Data and Generalizability** The primary limitation of this study is the specificity of the dataset, which consists of transcripts from fourth and fifth-grade mathematics classrooms in the United States. This narrow scope means our models lack proven generalizability to other grade levels, subject areas (e.g., literacy, science), or international school systems. While the underlying methods may hold broader potential, our specific findings are bound to the context of U.S. elementary mathematics education. Expanding the applicability of these models is contingent upon the availability of more varied public data.

**Task and Model Specificity** Our approach is limited by both the evaluation task and the model architecture. We focused on a subset of rating items from the MQI rubric, which may not fully represent the complexity of the universal task of instructional rating. Additionally, the inherent imperfections of observational rubrics, even for human experts, are a constraint on the ground truth data. Furthermore, our encoder models were custom-built for this task. While this specialized design allows a single model to score 25 distinct measures, it is not designed to generalize to new domains or contexts without substantial changes to its architecture or the introduction of new training data.

**Considerations for Practical Application and Deployment** Despite achieving high performance on several metrics, the models in their current state are not ready for high-stakes deployment. Substantial research and validation are necessary to ensure their reliability and to understand potential failure modes. Even when used with a human-in-the-loop, more work is needed to align the models' capabilities with the potential assumptions of end-users. Crucially, this study should not be interpreted as an endorsement for using general-purpose "GPT-style" large language models for similar evaluative tasks. The challenges inherent in this domain require specialized, carefully validated solutions rather than the application of general-purpose technologies.

## References

- Elena Aguilar. 2013. Developing a Work Plan: How Do I Determine What to Do? In *The art of coaching: effective strategies for school transformation*, pages 119–144. Jossey-Bass, A Wiley Brand, San Francisco.
- Sterling Alic, Dorottya Demszky, Zid Mancenido, Jing Liu, Heather Hill, and Dan Jurafsky. 2022. [Computationally Identifying Funneling and Focusing Questions in Classroom Discourse](#). *arXiv preprint*. ArXiv:2208.04715 [cs].
- Andrew Bacher-Hicks, Mark J. Chin, Thomas J. Kane, and Douglas O. Staiger. 2017. [An Evaluation of Bias in Three Measures of Teacher Quality: Value-Added, Classroom Observations, and Student Surveys](#).
- Andrew Bacher-Hicks, Mark J. Chin, Thomas J. Kane, and Douglas O. Staiger. 2019. [An experimental evaluation of three teacher quality measures: Value-added, classroom observations, and student surveys](#). *Economics of Education Review*, 73:101919.
- Paul Bambrick-Santoyo. 2016. *Get better faster: a 90-day plan for coaching new teachers*. Jossey-Bass, A Wiley Brand, San Francisco, CA.
- Paul Bambrick-Santoyo. 2018. *Leverage leadership 2.0: a practical guide to building exceptional schools*. Jossey-Bass, San Francisco, CA.
- Anthony J. Bishara and James B. Hittner. 2017. [Confidence intervals for correlations when data are not normal](#). *Behavior Research Methods*, 49(1):294–309.
- David Blazar. 2018. [Validating Teacher Effects on Students' Attitudes and Behaviors: Evidence from Random Assignment of Teachers to Students](#). *Education Finance and Policy*, 13(3):281–309.
- David Blazar and Cynthia Pollard. 2022. [Challenges and Tradeoffs of "Good" Teaching: The Pursuit of Multiple Educational Outcomes](#). Technical report, Annenberg Institute at Brown University. Publication Title: EdWorkingPapers.com.
- Robert L. Brennan. 2001. *Generalizability Theory*. Springer, New York, NY.
- Jodi M. Casabianca, Daniel F. McCaffrey, Drew H. Gitomer, Courtney A. Bell, Bridget K. Hamre, and Robert C. Pianta. 2013. [Effect of Observation Mode on Measures of Secondary Mathematics Teaching](#). *Educational and Psychological Measurement*, 73(5):757–783. Publisher: SAGE Publications Inc.
- Charalambos Y. Charalambous and Seán Delaney. 2019. [13 Mathematics Teaching Practices and Practice-Based Pedagogies](#). Brill. Section: International Handbook of Mathematics Teacher Education: Volume 1.
- Linda Darling-Hammond. 2014. [What Can PISA Tell Us about U.S. Education Policy?](#) *New England Journal of Public Policy*, 26(1).

- Linda Darling-Hammond, Lisa Flook, Channa Cook-Harvey, Brigid Barron, and David Osher. 2020. [Implications for educational practice of the science of learning and development](#). *Applied Developmental Science*, 24(2):97–140. Publisher: Routledge. eprint: <https://doi.org/10.1080/10888691.2018.1537791>.
- Dorottya Demszky and Heather Hill. 2022. [The NCTE Transcripts: A Dataset of Elementary Math Classroom Transcripts](#). Publisher: arXiv Version Number: 1.
- Dorottya Demszky and Jing Liu. 2023. [M-Powering Teachers: Natural Language Processing Powered Feedback Improves 1:1 Instruction and Student Outcomes](#). In *Proceedings of the Tenth ACM Conference on Learning @ Scale, L@S '23*, pages 59–69, New York, NY, USA. Association for Computing Machinery. Event-place: Copenhagen, Denmark.
- Dorottya Demszky, Jing Liu, Heather C. Hill, Shyamoli Sanghi, and Ariel Chung. 2023. [Improving Teachers' Questioning Quality through Automated Feedback: A Mixed-Methods Randomized Controlled Trial in Brick-and-Mortar Classrooms](#). Technical report, Annenberg Institute at Brown University. Publication Title: EdWorkingPapers.com.
- Dorottya Demszky, Jing Liu, Zid Mancenido, Julie Cohen, Heather Hill, Dan Jurafsky, and Tatsunori Hashimoto. 2021. [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, Online. Association for Computational Linguistics.
- Dorottya Demszky, Rose Wang, Sean Geraghty, and Carol Yu. 2024. [Does Feedback on Talk Time Increase Student Engagement? Evidence from a Randomized Controlled Trial on a Math Tutoring Platform](#). In *Proceedings of the 14th Learning Analytics and Knowledge Conference, LAK '24*, pages 632–644, New York, NY, USA. Association for Computing Machinery.
- Patrick J. Donnelly, Nathaniel Blanchard, Andrew M. Olney, Sean Kelly, Martin Nystrand, and Sidney K. D'Mello. 2017. [Words matter: automatic detection of teacher questions in live classroom discourse using linguistics, acoustics, and context](#). In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference, LAK '17*, pages 218–227, New York, NY, USA. Association for Computing Machinery.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2022. [SimCSE: Simple Contrastive Learning of Sentence Embeddings](#). *arXiv preprint*. ArXiv:2104.08821 [cs].
- Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, Lucy Campbell-Gillingham, Jonathan Uesato, Po-Sen Huang, Ramona Comanescu, Fan Yang, Abigail See, Sumanth Dathathri, Rory Greig, Charlie Chen, and 15 others. 2022. [Improving alignment of dialogue agents via targeted human judgements](#). *arXiv preprint*. ArXiv:2209.14375.
- Jason Grissom, Susanna Loeb, and Benjamin Master. 2013. [Effective Instructional Time Use for School Leaders: Longitudinal Evidence from Observations of Principals](#). *Educational Researcher*, 42(8)(42(8)):433.
- Zaretta Hammond. 2015. *Culturally responsive teaching and the brain: promoting authentic engagement and rigor among culturally and linguistically diverse students*. Corwin, a SAGE company, Thousand Oaks, California. OCLC: ocn889185083.
- Michael Hardy. 2024. ["All that Glitters": Approaches to Evaluations with Unreliable Model and Human Annotations](#). *arXiv preprint*. ArXiv:2411.15634 [cs].
- Michael Hardy. 2025. ["All that Glitters": Techniques for Evaluations with Unreliable Model and Human Annotations](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2250–2278, Albuquerque, New Mexico. Association for Computational Linguistics.
- 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. Publisher: Taylor & Francis, Ltd.
- Heather C. Hill, Charalambos Y. Charalambous, and Matthew A. Kraft. 2012. [When Rater Reliability Is Not Enough: Teacher Observation Systems and a Case for the Generalizability Study](#). *Educational Researcher*, 41(2):56–64. Publisher: American Educational Research Association.
- Andrew D. Ho and Thomas J. Kane. 2013. [The Reliability of Classroom Observations by School Personnel](#). Research Paper. MET Project. Technical report, Bill & Melinda Gates Foundation. Publication Title: Bill & Melinda Gates Foundation ERIC Number: ED540957.
- Harold Hotelling. 1936. [Relations Between Two Sets of Variates\\*](#). *Biometrika*, 28(3-4):321–377.
- Jennifer Jacobs, Karla Scornavacco, Charis Harty, Abhijit Suresh, Vivian Lai, and Tamara Sumner. 2022. [Promoting rich discussions in mathematics classrooms: Using personalized, automated feedback to support reflection and instructional change](#). *Teaching and Teacher Education*, 112:103631.



- Irina Jurenka, Markus Kunesch, Kevin R McKee, Daniel Gillick, Shaojian Zhu, Shubham Milind Phal, Katherine Hermann, Daniel Kasenberg, Avishkar Bhoopchand, Ankit Anand, Miruna Pislari, Stephanie Chan, Lisa Wang, Jennifer She, Parsa Mahmoudieh, Wei-Jen Ko, Andrea Huber, Brett Wiltshire, Gal Elidan, and 51 others. 2024. Towards Responsible Development of Generative AI for Education: An Evaluation-Driven Approach.
- Thomas Kane, Heather Hill, and Douglas Staiger. 2015. [National Center for Teacher Effectiveness Main Study: Version 4](#).
- 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. Technical report, Bill & Melinda Gates Foundation. Publication Title: Bill & Melinda Gates Foundation ERIC Number: ED540959.
- Thomas J. Kane and Douglas O. Staiger. 2012. [Gathering Feedback for Teaching: Combining High-Quality Observations with Student Surveys and Achievement Gains](#). Research Paper. MET Project. Technical report, Bill & Melinda Gates Foundation. Publication Title: Bill & Melinda Gates Foundation ERIC Number: ED540960.
- 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. Publisher: American Educational Research Association.
- Maurice George Kendall. 1938. [A NEW MEASURE OF RANK CORRELATION](#). *Biometrika*, 30(1-2):81–93.
- Maurice George Kendall. 1945. [THE TREATMENT OF TIES IN RANKING PROBLEMS](#). *Biometrika*, 33(3):239–251.
- Doug Lemov. 2021. *Teach like a champion 3.0: 63 techniques that put students on the path to college*, third edition edition. Jossey-Bass, a Wiley imprint, Hoboken, NJ.
- Doug Lemov and Norman Atkins. 2015. *Teach like a champion 2.0: 62 techniques that put students on the path to college*, second edition edition. Jossey-Bass, San Francisco, CA.
- Zehan Li, Xin Zhang, Yan Zhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. [Towards General Text Embeddings with Multi-stage Contrastive Learning](#). *arXiv preprint*. ArXiv:2308.03281 [cs].
- Peter Liljedahl, Tracy Johnston Zager, and Laura Wheeler. 2021. *Building thinking classrooms in mathematics: 14 teaching practices for enhancing learning: Grades K-12*. Corwin Mathematics. Corwin, Thousand Oaks, California London New Delhi Singapore.
- Jing Liu and Julie Cohen. 2021. [Measuring Teaching Practices at Scale: A Novel Application of Text-as-Data Methods](#). *Educational Evaluation and Policy Analysis*, 43(4):587–614. Publisher: American Educational Research Association.
- 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 [cs].
- Robert C. Pianta, Jay Belsky, Nathan Vandergrift, Renate Houts, and Fred J. Morrison. 2008. [Classroom Effects on Children's Achievement Trajectories in Elementary School](#). *American Educational Research Journal*, 45(2):365–397. Publisher: American Educational Research Association.
- Robert C. Pianta and Bridget K. Hamre. 2009. [Conceptualization, Measurement, and Improvement of Classroom Processes: Standardized Observation Can Leverage Capacity](#). *Educational Researcher*, 38(2):109–119. Publisher: American Educational Research Association.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks](#). *arXiv preprint*. ArXiv:1908.10084 [cs].
- Borhan Samei, Andrew M. Olney, Sean Kelly, Martin Nystrand, Sidney D'Mello, Nathan Blanchard, Xiaoyi Sun, Marcy Glaus, and Art Graesser. 2014. [Domain Independent Assessment of Dialogic Properties of Classroom Discourse](#). Technical report. Publication Title: Grantee Submission ERIC Number: ED566380.
- Jon Saphier, Mary Ann Haley-Speca, and Robert Gower. 2008. *The skillful teacher: building your teaching skills*, 6th ed edition. Research for Better Teaching, Acton, Mass.
- Daniel L. Schwartz, Jessica M. Tsang, and Kristen P. Blair. 2016. *The ABCs of how we learn: 26 scientifically proven approaches, how they work, and when to use them*, first edition edition. Norton books in education. W.W. Norton & Company, New York.
- 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 [cs].
- Anais Tack, Ekaterina Kochmar, Zheng Yuan, Serge Bibauw, and Chris Piech. 2023. [The BEA 2023 Shared Task on Generating AI Teacher Responses in Educational Dialogues](#). *arXiv preprint*. ArXiv:2306.06941.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and

- Furu Wei. 2024a. [Text Embeddings by Weakly-Supervised Contrastive Pre-training](#). *arXiv preprint*. ArXiv:2212.03533 [cs].
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024b. [Multilingual E5 Text Embeddings: A Technical Report](#). *arXiv preprint*. ArXiv:2402.05672 [cs].
- Rose Wang and Dorottya Demszky. 2023. [Is ChatGPT a Good Teacher Coach? Measuring Zero-Shot Performance For Scoring and Providing Actionable Insights on Classroom Instruction](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 626–667, Toronto, Canada. Association for Computational Linguistics.
- Jacob Whitehill and Jennifer LoCasale-Crouch. 2024. [Automated Evaluation of Classroom Instructional Support with LLMs and BoWs: Connecting Global Predictions to Specific Feedback](#). *arXiv preprint*. ArXiv:2310.01132 [cs].
- Grover J. Whitehurst, Matthew M. Chingos, and Katharine M. Lindquist. 2014. [Evaluating Teachers with Classroom Observations: Lessons Learned in Four Districts](#). Technical report, Brookings Institution. Publication Title: Brookings Institution ERIC Number: ED553815.
- Paiheng Xu, Jing Liu, Nathan Jones, Julie Cohen, and Wei Ai. 2024. [The Promises and Pitfalls of Using Language Models to Measure Instruction Quality in Education](#). *arXiv preprint*. ArXiv:2404.02444 [cs].
- Grace Yoon, Raymond J Carroll, and Irina Gaynanova. 2020. [Sparse semiparametric canonical correlation analysis for data of mixed types](#). *Biometrika*, 107(3):609–625.

## **A Item Information**

## **B Figures**

### **B.1 Human Expert Score Distributions**

### **B.2 Test Set Distributions**




Figure 4: **Human Expert Score Distributions.** These are the score distributions from human experts. The distinct rating patterns highlight the underlying qualitative differences in the constructs being rated. Previous studies have focused on a limited range of items (bolded, (Wang and Demszky, 2023))

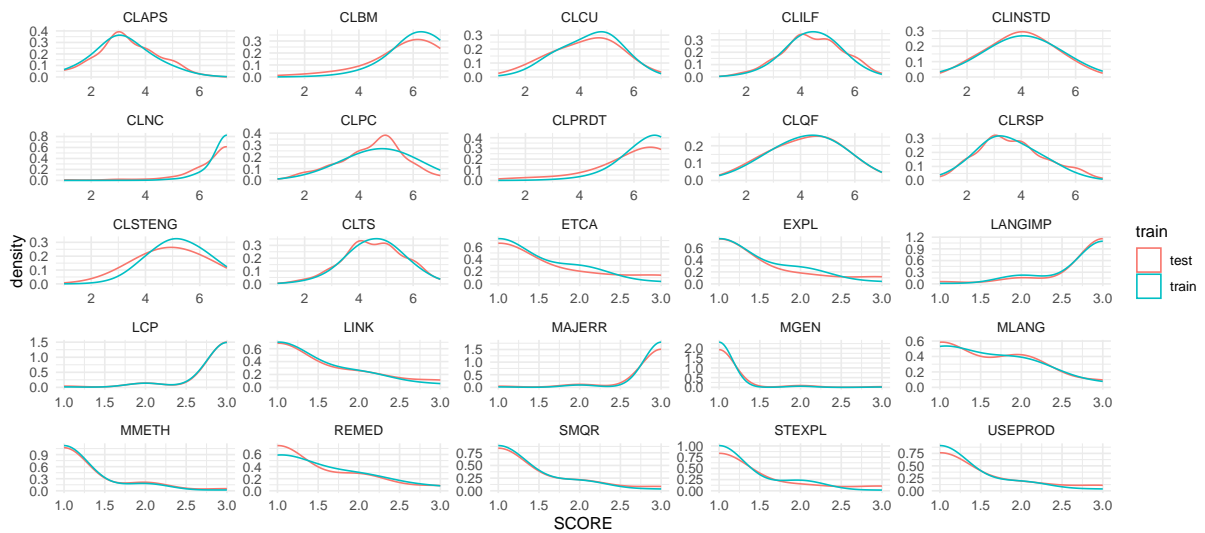


Figure 5: **Held-out Test Set Distributions.** These are comparative score distributions from human experts for the items in the held-out test set and the remaining sample. No differences were statistically significant.

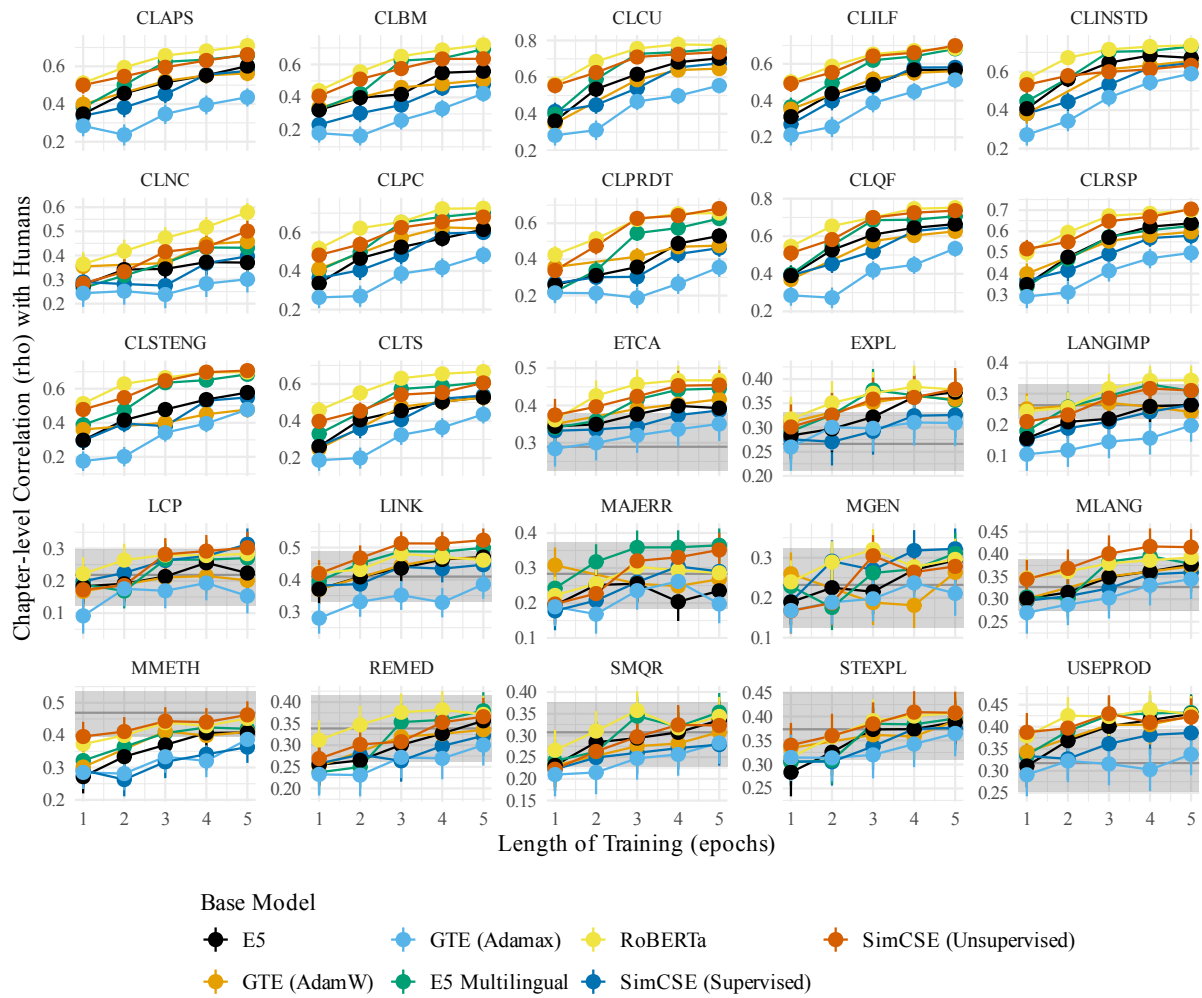


Figure 6: Chapter-level by Item-level Correlation with human experts across training epochs




Figure 7: The MQI instrument had at least two human raters per lesson, and the mean and interquartile range of all 63 human MQI raters correlated across the other raters are represented by the gray line and shaded region in the figure. When looking at only the specific chapter from the lesson, however, models only perform in the top half of human raters. This discrepancy may be an artifact of the windowing structure used in the training. The multilevel framework decomposes the correlation into within- and between-cluster components:  $\rho_{\text{partial}} = \text{Corr}(\text{rank}(\tilde{R}_{ij}^{(1)}), \text{rank}(\tilde{R}_{ij}^{(2)}) \mid \mathbf{X}_{ij})$  where  $\tilde{R}_{ij}^{(k)}$  represents residualized ratings after partialling out covariates  $\mathbf{X}_{ij}$  from the multilevel model:  $R_{ij}^{(k)} = \beta^{(k)} \mathbf{X}_{ij} + u_j^{(k)} + \epsilon_{ij}^{(k)}$  with random intercepts  $u_j^{(k)} \sim N(0, \tau^2)$  for items and residuals  $\epsilon_{ij}^{(k)} \sim N(0, \sigma^2)$ . The multilevel partial Spearman’s correlation accounts for the hierarchical structure while providing a robust, rank-based measure of association that generalizes beyond the specific items sampled, with inference based on the random effects distribution of items.

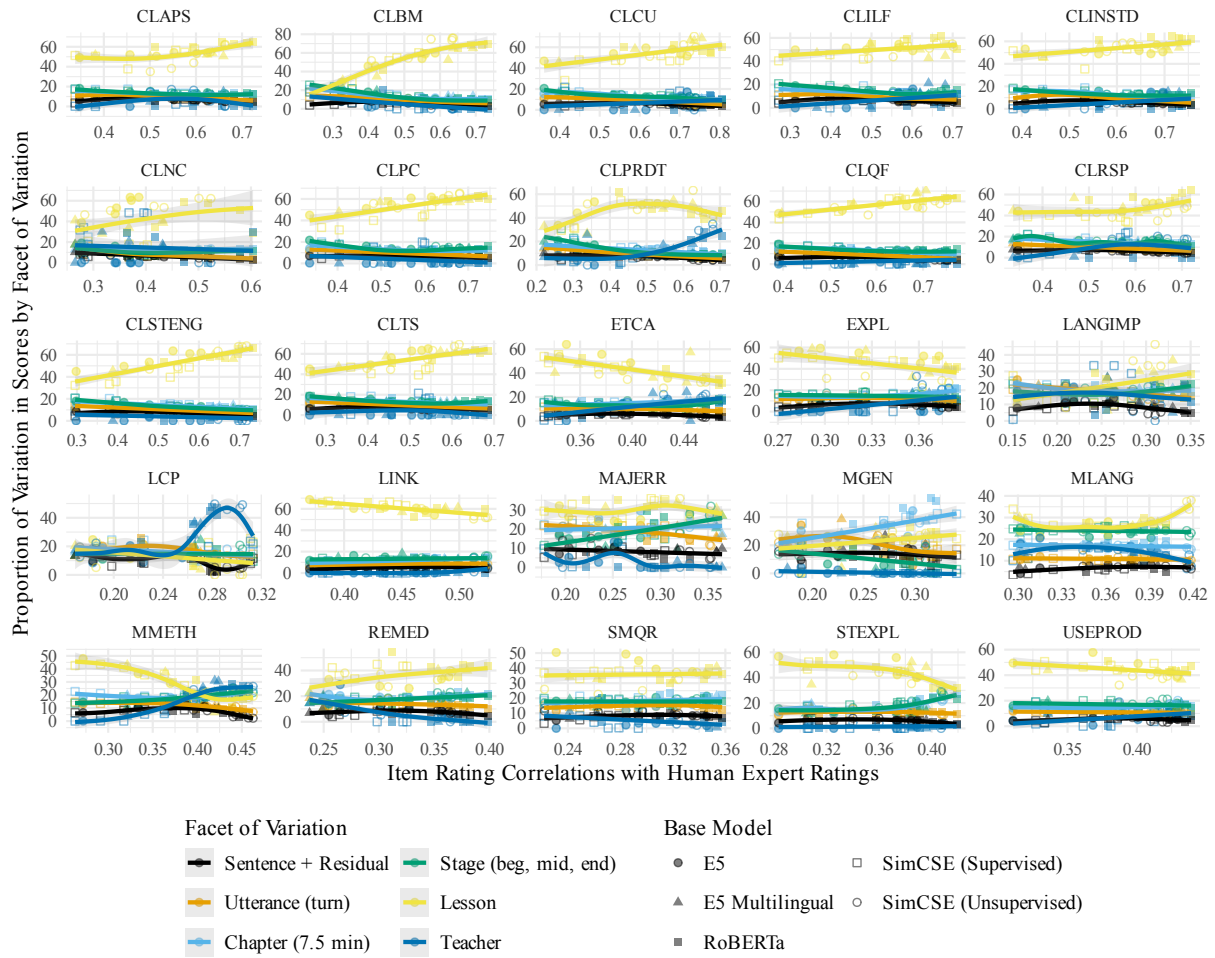


Figure 8: Proportion of variation explained as related to a model’s alignment to human expert ratings disaggregated by item using generalizability theory. Random effects were calculated using the lme4 package in R: SCORE (1|NCTETID) + (1|OBSID) + (1|OBS\_CHAPS) + (1|OBS\_CHAP) + (1|OBS\_CHAP\_idx)

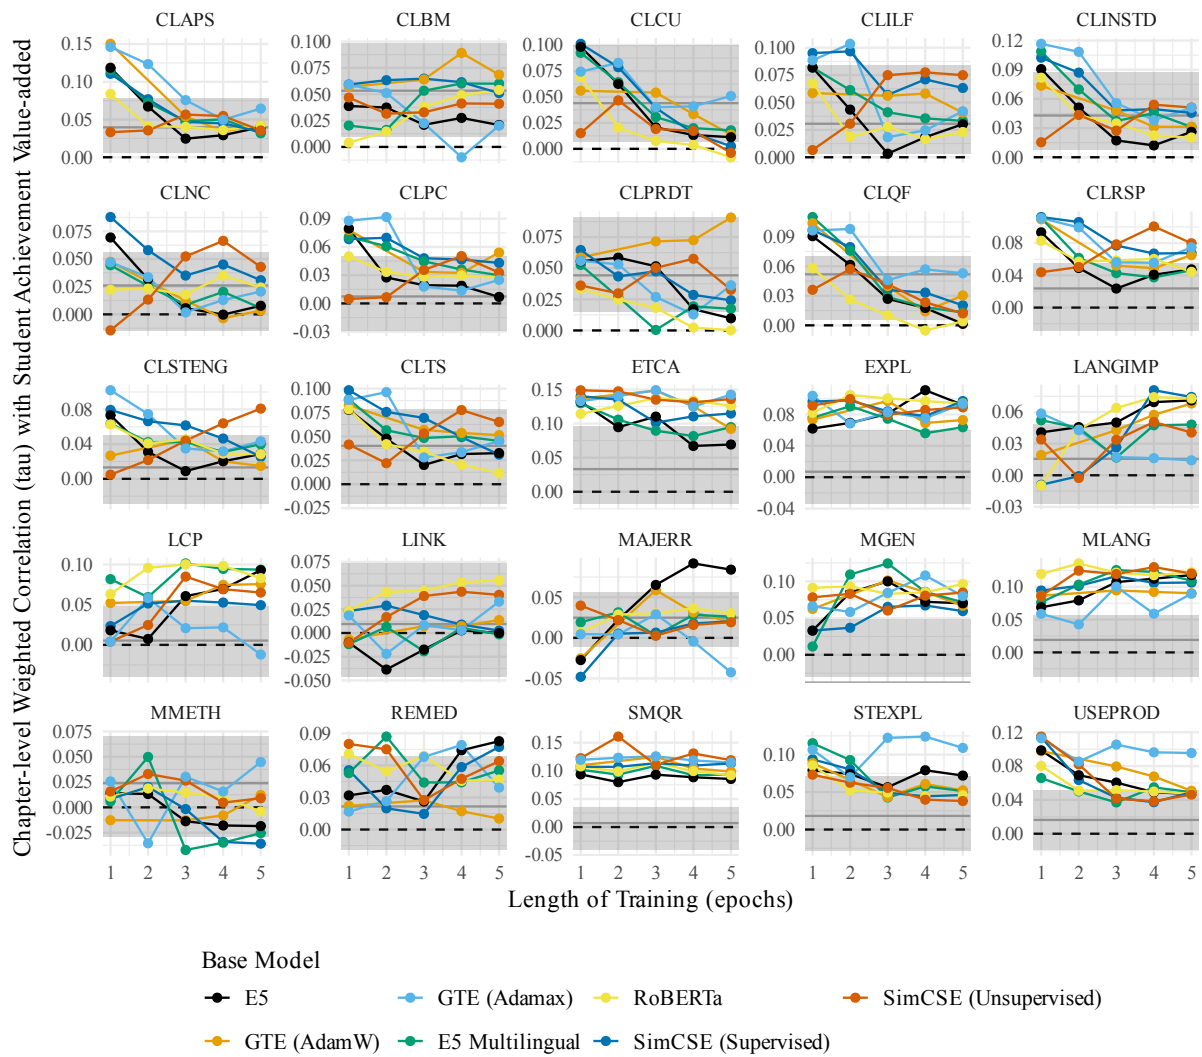


Figure 9: Kendall tau correlations on the stacked VAM outcomes disaggregated at the item level as a function of training epoch. The mean and interquartile range of all human raters evaluated in the same manner are represented by the gray line and shaded region in the figure. Weights are used to account for number of observations per unit.




Figure 10: Lesson-level canonical correlations as a function of correlation with human ratings.



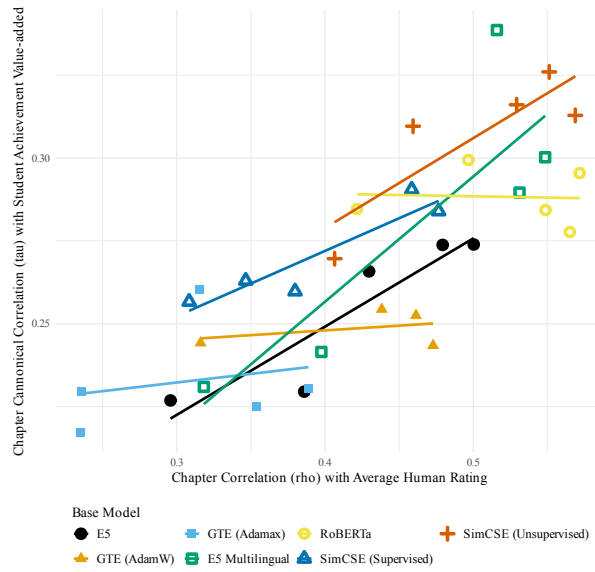


Figure 11: Chapter-level canonical correlations using methods from Yoon et al. as a function of correlation with human ratings.




Figure 12: Lesson-level canonical correlations using methods from Yoon et al. as a function of correlation with human ratings.

<b>Abbreviation</b>	<b>Item</b>	<b>Item Description</b>
<b>MQI Instrument</b>		
ETCA	<i>Enacted Task Cognitive Activation</i>	Task cognitive demand, such as drawing connections among different representations, concepts, or solution methods; identifying and explaining patterns.
<b>EXPL</b>	<i>Teacher Explanations</i>	Teacher explanations that give meaning to ideas, procedures, steps, or solution methods.
<b>LANGIMP</b> †	<i>Imprecision in Language or Notation</i>	Imprecision in language or notation, with regard to mathematical symbols and technical or general mathematical language.
LCP†	<i>Lack of Clarity in Presentation of Mathematical Content</i>	Lack of clarity in teachers' launching of tasks or presentation of the content.
LINK	<i>Linking and Connections</i>	Linking and connections of mathematical representations, ideas, and procedures.
MAJERR†	<i>Major Mathematical Errors</i>	Major mathematical errors, such as solving problems incorrectly, defining terms incorrectly, forgetting a key condition in a definition, equating two non-identical mathematical terms.
MGEN	<i>Developing Mathematical Generalizations</i>	Developing generalizations based on multiple examples.
MLANG	<i>Mathematical Language</i>	Mathematical language is dense and precise and is used fluently and consistently.
MMETH	<i>Multiple Procedures or Solution Methods</i>	Multiple procedures or solution methods for a single problem.
<b>REMED</b>	<i>Remediation of Student Errors and Difficulties</i>	Remediation of student errors and difficulties addressed in a substantive manner.
<b>SMQR</b>	<i>Student Mathematical Questioning and Reasoning</i>	Student mathematical questioning and reasoning, such as posing mathematically motivated questions, offering mathematical claims or counterclaims.
STEXPL	<i>Students Provide Explanations</i>	Student explanations that give meaning to ideas, procedures, steps, or solution methods.
USEPROD	<i>Responding to Student Mathematical Productions</i>	Responding to student mathematical productions in instruction, such as appropriately identifying mathematical insight in specific student questions, comments, or work; building instruction on student ideas or methods.
<b>CLASS Instrument</b>		
<b>CLPC</b>	<i>Classroom Positive Climate</i>	
CLNC†	<i>Classroom Negative Climate</i>	
CLTS	<i>Teacher Sensitivity</i>	
CLRSP	<i>Regard for Student Perspective</i>	
<b>CLBM</b>	<i>Behavior Management</i>	
CLPRDT	<i>Productivity</i>	
CLILF	<i>Instructional Learning Formats</i>	
CLCU	<i>Content Understanding</i>	
CLAPS	<i>Applied Problem Solving</i>	
CLQF	<i>Quality of Feedback</i>	
<b>CLINSTD</b>	<i>Instructional Dialogue</i>	
CLSTENG	<i>Student Engagement</i>	

Table 1: CLASS and MQI item descriptions and corresponding abbreviations. †denotes items that are reverse coded due to being negatively worded with respect to the construct of teacher ability. Bolded items are those evaluated by the **GPT** family of raters and reported by Wang and Demszky. Each member of the Human and Encoder families of raters evaluated all 25 items.