

Utility-Preserving De-Identification for Math Tutoring: Investigating Numeric Ambiguity in the MathEd-PII Benchmark Dataset

Zhuqian Zhou
Cornell University
zz968@cornell.edu

Kirk Vanacore
Cornell University
kpv27@cornell.edu

Bakhtawar Ahtisham
Cornell University
ba453@cornell.edu

Jinsook Lee
Cornell University
jl3369@cornell.edu

Doug Pietrzak
Fresh Cognate
doug@freshcognate.com

Daryl Hedley
Fresh Cognate
daryl@freshcognate.com

Jorge Dias
Fresh Cognate
jorge@freshcognate.com

Chris Shaw
UPchieve
chris.shaw@upchieve.org

Ruth Schäfer
Saga Education
DrRuthSchafer@gmail.com

René F. Kizilcec
Cornell University
kizilcec@cornell.edu

ABSTRACT

Large-scale sharing of dialogue data is key to advancing the science of teaching and learning, yet rigorous de-identification remains a major barrier. In mathematics tutoring transcripts, numeric expressions frequently resemble structured identifiers (e.g., dates or IDs), leading generic Personally Identifiable Information (PII) detection systems to over-redact core instructional content and reduce data utility. This work asks how to detect PII while preserving educational utility, focusing on this “numeric ambiguity” problem. We introduce MathEd-PII, the first benchmark dataset for PII detection in math tutoring dialogues, built with human-in-the-loop LLM annotation. Using density-based segmentation, we show that false PII redactions cluster in math-dense regions, confirming numeric ambiguity as a key failure mode. We then compare four detection strategies: a Presidio baseline and three LLM-based approaches with basic, math-aware, and segment-aware prompting. Domain-aware prompting, including both math-aware (F1: 0.802) and segment-aware versions (F1: 0.821), substantially outperforms the baseline (F1: 0.379) while reducing numeric false positives, demonstrating that de-identification must incorporate domain context to preserve analytic utility. This work provides a new benchmark and evidence that utility-preserving de-identification for tutoring data requires domain-aware modeling.

Keywords

De-identification, PII Detection, Math Tutoring, Large Language Models, Privacy, Segmentation

Zhuqian Zhou, Kirk Vanacore, Bakhtawar Ahtisham, Doug Pietrzak, Daryl Hedley, Jorge Dias, Jinsook Lee, Rene Kizilcec, Chris Shaw, and Ruth Schäfer. Utility-Preserving De-Identification for Math Tutoring: Investigating Numeric Ambiguity in the MathEd-PII Benchmark Dataset. In Anthony Botelho, Maria Mercedes T. Rodrigo, Adish Singla, Hiroaki Ogata, Hyejeong So, and Young Hoan Cho (eds.) Proceedings of the 19th International Conference on Educational Data Mining, Seoul, Republic of Korea, June, 2026, pp. 365–375. International Educational Data Mining Society (2026).

© 2026 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.

<https://doi.org/10.5281/zenodo.21039941>

1. INTRODUCTION

The Educational Data Mining (EDM) community has long relied on large-scale open source data sets from digital learning platforms [14, 18]. Many of these data sources can be easily de-identified because they are composed of action logs that, for the most part, do not typically contain personally identifiable information (PII). As more and more educational technologies provide dialogue-based tutoring at scale, either with humans or with artificial intelligence (AI) chatbots [3, 8, 22, 26], access to large corpora of tutoring conversations has become essential. However, sharing such data requires rigorous de-identification to remove direct identifiers (e.g., names and email addresses) and quasi-identifiers (e.g., gender, ages, locations, and school names) that could enable re-identification of specific individuals.

Recent work has begun to develop tools and workflows for de-identifying educational text [9, 10, 24, 25, 31, 32]. Yet specialized domains like mathematics tutoring introduce a distinct and underexplored challenge. In math dialogue, numbers are the primary medium of exchange. A teacher might ask, “Are $4/12$ and $2/6$ the same?” (fractions), or a student might calculate “ $500 - 70 - 2000$ ” (subtraction). Standard PII detectors frequently misidentify these mathematical expressions as DATE or US_SSN due to *numeric ambiguity*, the overlap between the surface forms of mathematical content and sensitive identifiers. Over-identifying these math-related spans results in low precision, stripping the data of its core educational content and rendering it useless for downstream tasks.

A second barrier is the lack of domain-specific benchmarks. To date, no publicly available dataset exists for evaluating PII detection in math tutoring contexts, limiting the ability to rigorously study and improve de-identification methods for this increasingly important data modality.

To overcome these barriers, this work addresses the following research question (RQ): How can we accurately detect PII in math tutoring transcripts while preserving the educational utility of the data? We operationalize this question

through three sub-questions that structure the remainder of the paper:

RQ1 (Failure analysis): To what extent does numeric ambiguity in mathematical discourse lead to false positive PII detection?

RQ2 (Dataset construction): How can a reliable benchmark for evaluating PII detection be created when original unredacted data cannot be shared?

RQ3 (Method evaluation): Do math-aware prompting strategies improve utility-preserving PII detection compared to domain-agnostic baselines?

To answer these questions, we begin with a large corpus of math tutoring transcripts that had already been redacted for PII by the data provider. The original unredacted data are unavailable because the data provider’s internal data safety regulation requires responsible removal of PII before any data sharing can happen. Therefore, we developed a human-in-the-loop large language model (LLM) workflow to audit and repair PII redactions done by the data provider before they shared the data with us and to generate privacy-preserving surrogates. This process yields **MathEd-PII**¹, a benchmark dataset of 1,000 tutoring sessions (115,620 messages; 769,628 tokens) for evaluating PII detection in math tutoring dialogue.

We find that numeric expressions are a major source of false PII detections in math tutoring dialogue. Using a segmentation analysis to examine how detection errors differ between math-dense and conversational regions of transcripts, we observe 55.5% of false redactions to occur in math-dense segments that cover only 36.6% of the corpus, whereas true PII appears substantially less often in these regions (22.2%). We evaluate four PII detection strategies: (i) a Presidio-based baseline, (ii) LLMs with basic prompting, (iii) LLMs with math-aware prompting, and (iv) LLMs with segment-aware prompting. Math-aware prompting substantially improves performance over the baseline (F1: 0.821 vs. 0.379), primarily by reducing numeric false positives. These findings demonstrate that de-identification for tutoring dialogue cannot be treated as a domain-agnostic preprocessing step; preserving analytic utility requires incorporating mathematical discourse context.

This paper makes three contributions aligned with the RQs above. We identify numeric ambiguity as a central failure mode in de-identifying math tutoring dialogue (RQ1). We then introduce MathEd-PII, the first benchmark dataset for PII detection in math tutoring transcripts (RQ2). Finally, we show that math-aware prompting substantially improves utility-preserving de-identification (RQ3).

2. RELATED WORK

2.1 Policy Context

De-identification is a prerequisite for sharing educational interaction data at scale, but what counts as “sufficient” de-identification depends on the legal regime and the assumed adversary model. In the United States, the Children’s Online Privacy Protection Act (COPPA) and the Family Educational Rights and Privacy Act (FERPA) govern education

records and only permits disclosure of de-identified information when there is no reasonable basis to believe the remaining information can be used to identify an individual [4, 29].

In parallel, the Health Insurance Portability and Accountability Act (HIPAA) provides two common pathways for de-identification of health information—*Safe Harbor* (removal of 18 enumerated identifiers, e.g., names, geographic subdivisions smaller than a state, dates directly related to an individual, and unique identification numbers) and *Expert Determination* (a documented, risk-based assessment) [30]. Although specific to health data, the 18 identifiers are widely used as a practical reference beyond that domain because they offer “the promise of a straightforward application of rules, a repeatable process, and a known result: a dataset that is legally de-identified” [6, p. 24]. Despite this apparent clarity, the National Institute of Standards and Technology (NIST) emphasizes that de-identification is inherently dynamic, framing it along an “identifiability spectrum” and cataloging transformation families such as suppression, generalization, perturbation, masking, and synthesis. Stronger privacy protections, it notes, typically come at the expense of analytic fidelity [6, 7].

2.2 PII Detection in Unstructured Text

Most practical pipelines for unstructured text de-identification begin with *entity detection* (identifying spans that correspond to PII) and then apply an *anonymization action* (e.g., redaction, masking, or surrogate replacement). Historically, detection approaches fall into three broad categories—rule-based, model-based, and LLM-based methods.

Rule-based PII detection systems rely on hand-engineered patterns and resources, including regular expressions for structured identifiers (e.g., emails, phone numbers, and dates), curated lexicons for names and locations, and deterministic context heuristics [21]. These methods can be highly precise for well-formed identifiers but brittle under conversational noise (misspellings, fragments, idiosyncratic formatting), and they require continual maintenance as domains shift [19].

A complementary line of work frames PII detection as a named entity recognition (NER) task, in which statistical or neural models predict entity spans over tokens or characters based on contextual cues. Contemporary NLP toolkits, most notably spaCy, provide pretrained and fine-tunable NER models that are commonly used as baseline PII detectors or adapted to domain-specific corpora [12]. In practice, many production systems combine *rule-based* and *model-based* paradigms to exploit their complementary strengths. Microsoft Presidio exemplifies this hybrid design by integrating rule-based recognizers for high-precision patterns with neural-network-based entity detectors within a unified anonymization framework [16].

More recently, LLMs have been applied to de-identification both via prompting and fine-tuning [24, 25, 31]. In prompting-based setups, the model is typically instructed (often with a schema) to tag or extract PII spans and entities from raw text—either by returning structured outputs (e.g., JSON with character offsets and labels) or by

¹<https://huggingface.co/datasets/NationalTutoringObservatory/MathEd-PII>

emitting an annotated/redacted version of the transcript—whereas fine-tuning-based setups train the model on labeled examples to produce consistent token-level tags or span predictions under a fixed output format [24].

2.3 Augmented and Synthetic data for PII Detection and Evaluation

A practical challenge in PII detection research is the scarcity of richly annotated, shareable corpora, as releasing PII labels alongside text can increase re-identification risk. To address this, two complementary strategies are increasingly common: data augmentation and synthetic data generation.

Data augmentation for PII detection starts from de-identified or manually redacted text and inserts synthetic PII spans, often through surrogate replacement schemes such as “hiding in plain sight,” which preserve discourse coherence while reducing disclosure risk [2, 27]. In parallel, fully synthetic datasets have been constructed using template- or rule-driven generators and, more recently, LLMs conditioned on schema constraints (e.g., entity types, formats, and conversational style) [17, 23, 28]. These approaches offer controllability and improved coverage of rare entity types, but their effectiveness depends on the realism of the generated text: synthetic data that fails to capture domain-specific language, error patterns, or contextual ambiguity may yield overly optimistic performance estimates and limited transfer to real-world settings [23, 28].

2.4 De-Identification of Educational Data

Early work on de-identification in educational text focused on providing practical tooling and datasets to support privacy-preserving analysis of student writing and interaction data. Holmes et al. [9, 10] introduced PILO, an open-source system for labeling and obfuscating PII in educational corpora, and examined de-identification challenges in student-authored texts across digitally mediated learning environments. Building on this foundation, Singhal et al. [25] conducted one of the first systematic evaluations of LLMs for educational data de-identification, using online course discussion forums as a testbed and demonstrating that LLMs substantially improve recall but often over-redact, motivating the need for utility-aware evaluation. Subsequent work by Zambrano et al. [31] expanded this analysis with larger-scale experiments and error analyses, further documenting precision–utility trade-offs in LLM-based anonymization of student forum posts. More recently, Shen et al. [24] proposed enhanced LLM-based pipelines that combine prompting and verification strategies to improve robustness across educational datasets, while Zent et al. [32] introduced an anonymization framework tailored to question-anchored tutoring dialogues, emphasizing deployability and domain adaptation. Together, this line of work reflects a rapid shift from rule-based and toolkit-oriented approaches toward LLM-centric methods, alongside growing recognition that educational domains pose distinct challenges for balancing privacy protection and analytic utility as it is conversational and tightly interleaves task content with incidental personal references.

Math tutoring transcripts present an additional challenge: beyond conversational noise, they contain dense, non-

sensitive numerals and operators whose surface forms overlap with common PII categories such as AGE, DATE, and ID. Over-redacting this math-related content would severely reduce the downstream educational and research-related utility of the dataset. However, numerical false positives in PII detection within mathematical contexts have received little attention. A major reason is the scarcity of labeled data: there are only two publicly available educational datasets with PII annotations, one consisting of student-authored articles submitted to a massive open online course (MOOC) on critical thinking through design² [11] and the other containing teacher–student chatroom messages for English language learning³ [1]. To our knowledge, no public datasets exist for PII detection in math education contexts.

3. CURRENT RESEARCH OVERVIEW

To answer the research questions outlined above, we follow a three-phase research workflow moving from dataset construction to analytical validation and finally to evaluation of math-aware de-identification strategies.

Phase 1: Dataset Preparation involves constructing MathEd-PII, a benchmark dataset for PII detection in math tutoring context. The construction includes identifying original over-redaction and establishing a clean ground truth through human-in-the-loop LLM-powered PII surrogation.

Phase 2: Math Segmentation focuses on developing a density-based module to isolate math-heavy regions, serving as an analytical lens to validate the hypothesis of numeric ambiguity that false PII redactions are structurally concentrated in mathematical content while true redactions are not.

Phase 3: PII Detection and Evaluation implements and compares four PII detection strategies—a Presidio-based baseline, LLMs with a basic prompt, a general math-domain-aware prompt, and a specific math-segment-aware prompt—evaluating Precision (P), Recall (R), and F1 overall, as well as stratified by PII types and by math versus non-math segments.

4. PHASE 1: DATASET PREPARATION

Reference datasets for PII detection in math education are currently lacking. To enable rigorous evaluation, we constructed a benchmark dataset, **MathEd-PII**, from a PII-redacted large corpus.

4.1 Source Corpus

Our source corpus comprises 1,000 math tutoring sessions (115,620 messages; 769,628 tokens) from a U.S.-based tutoring platform. Among these, 4,648 (4%) messages spread in 786 sessions feature 5,265 PII redactions applied by the platform’s in-house PII detection system—Microsoft Presidio, which, as introduced in 2.2, is an industry-standard open-source software development kit (SDK) for PII detection by integrating rule-based recognizers for high-precision patterns with neural-network-based entity detectors within a unified anonymization framework [16] with customizable

²Available at <https://www.kaggle.com/datasets/langdonholmes/cleaned-repository-of-annotated-pii/data>

³Available via request at forms.gle/pKc48WMhnySC8zDk9

rules. The provider redacts 12 types of PII, such as PERSON, LOCATION, US_DRIVER_LICENSE, etc. The full list of PII types can be found in Table 6.

Initial inspection suggests that several redaction types are not consistently reliable indicators of true PII in math tutoring dialogue. In particular, some types that resemble structured numeric identifiers (e.g., COURSE_NUMBER and US_DRIVER_LICENSE) are frequently triggered by math content rather than identity disclosure (see an example of incorrect redactions and inferred original text⁴ in Table 1). Other types (notably LOCATION, NRP⁵, and URL) appear mixed, containing both correct redaction and substantial over-redaction.

Table 1: Incorrect PII Redactions in the Source Corpus and Inferred Original Text

Redacted Message	Inferred Original Message
9351 / 7 = 1335 <US_DRIVER_LICENSE>	9351 / 7 = 1335.8571
go from radians to degrees it would <COURSE_NUMBER>/pi	go from radians to degrees it would 180/pi
so it would be 1.15x= 1.15x=<PHONE_NUMBER>368000	so it would be 1.15x= 1.15x=

4.2 Method: Human-In-The-Loop PII Redaction Evaluation and Surrogate Generation

To evaluate the extent of incorrect redactions and construct a reliable benchmark for PII detection from the current dataset, we employed a human-in-the-loop LLM workflow to assess the quality of existing PII labels and generate plausible surrogates, thereby restoring transcript completeness for downstream investigation. This design choice is motivated by three practical constraints. First, the original un-redacted transcripts are unavailable due to privacy and data-protection requirements, precluding direct comparison against ground-truth PII. Second, manually reviewing and reconstructing every redaction at scale would be prohibitively costly and time-consuming. Third, LLMs are well-suited for this task because they can leverage rich contextual information to infer the most plausible lexical realizations consistent with surrounding discourse, enabling systematic approximation of redacted content without reintroducing real personal identifiers. Human oversight is incorporated throughout the workflow to validate model outputs, refine prompts, and ensure that generated surrogates remain contextually coherent while preserving privacy. This workflow consists of the following three steps and is illustrated in Figure 1.

First, **Initial LLM-Based Assessment**: We first drafted an initial prompt that instructed the LLM to evaluate messages from a randomly sampled set of 100 transcripts, using a

⁴Original un-redacted text is not available due to the data provider’s privacy compliance. The original text here is inferred by an LLM based on the context. Please see Section 4.2 for more details.

⁵A person’s Nationality, Religious or Political group.

context window of at least three utterances preceding and following each target message. For each message, the LLM was prompted to identify any redacted PII or PII missed by the data provider’s internal redaction process, determine its PII type, and classify the redaction as “PII”, “Not PII”, or “Uncertain” based on contextual evidence. When PII was identified, the LLM was further instructed to generate context-appropriate surrogate text that preserved semantic coherence while remaining substantially different from any original PII, when available.

Second, **Human-In-The-Loop Validation and Refinement**:

Three researchers—the first three authors of this paper—then manually reviewed the LLM outputs. Each researcher was assigned one-third of the outputs and down-voted instances in which any component of the output was deemed incorrect, including the identified PII type, quality classification, or the generated surrogate. Based on these annotations, the prompt was revised iteratively to address observed failure modes. This process continued until at least 95% of previously down-voted issues were successfully resolved by the updated prompt.

Third, **Pipeline Execution at Scale**: The final, validated prompt⁶ for redaction evaluation and surrogate generation was subsequently applied to the full 1,000-transcript dataset to produce the benchmark used in downstream analyses.

All three steps were conducted using an open-sourced annotation web application, Sandpiper⁷, optimized for stable, human-in-the-loop annotation iteration. The system provides a robust, interactive user interface that allows researchers to directly review, up-vote, or down-vote AI-generated annotations, enabling efficient human verification and correction at scale. Its modular architecture and worker-based execution model ensure stable handling of large annotation workloads and repeated LLM inference, supporting consistent evaluation across iterations.

4.3 Source Corpus Evaluation Results

Applying the procedures described above, we discovered both recall and precision issues of the data provider’s internal PII redaction presented in the source corpus as described below.

The Recall Issue: The human-in-the-loop LLM workflow discovered 201 messages containing PII that the data provider’s system failed to redact, spanning PERSON (159), GRADE_LEVEL (23), AGE (8), DATE (4), COURSE_NUMBER (3), NRP⁸ (3), SCHOOL (2), and URL (1). Notably, DATE and AGE were newly introduced PII types under the Safe Harbor framework (see Section 2.1) and were not included among the original dataset’s 12 PII categories. Still, this demonstrates a substantial recall issue: a widely used approach to de-identification, Microsoft Presidio, failed to identify 159 names in tutoring transcripts. Table 2 presents simulated examples of missed names. These examples were generated from failure patterns observed in the original dataset and do not refer to any real

⁶see 1 in https://github.com/ZhuqianZhou/MathEd-PII/blob/69012c0dbde3b3976f36533e2e44f5a5bfdfd718/EDM2026_paper_183_Appendix.pdf

⁷<https://sandpiperresearch.org/>

⁸see foot note 5

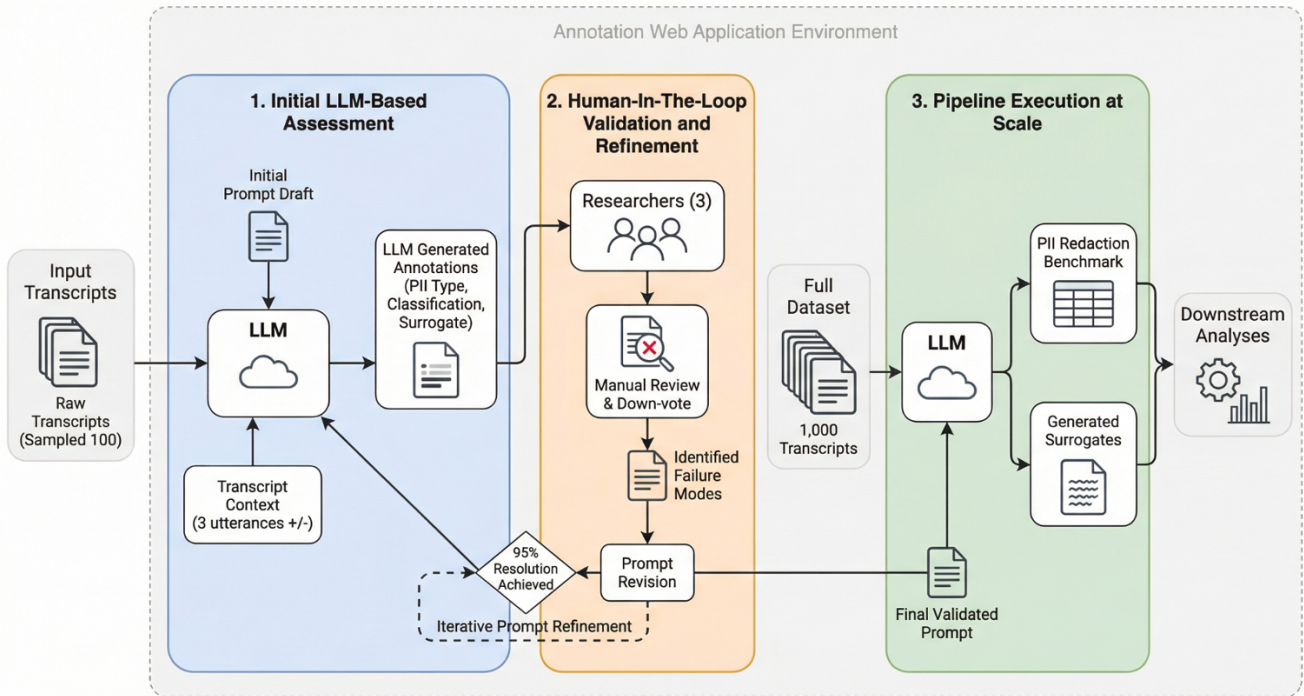


Figure 1: The Three-Step Human-In-The-Loop MathEd-PII Benchmark Dataset Preparation Workflow.

people. Readers can replicate the same failure using Presidio’s default Analyzer Engine with customized rules.

Table 2: Simulated Examples of PERSON that Presidio Failed to Detect and Similar Cases that Presidio Successfully Detected

Unsuccessful Detection	Successful Detection
hey olivia	hey Olivia
Hi Alex	Hi Alex,
I am Lam	i am Lam

The Precision Issue: Of the 4,648 originally redacted messages, 3,015 (64.87%) contained redactions evaluated as “Not PII”, indicating low precision in their internal system. Several PII types contained predominantly false positives, with the following four types receiving more than 90% “Not PII” flags—US_DRIVER_LICENSE (99.36%), COURSE_NUMBER (96.65%), PHONE_NUMBER (93.33%), and NRP (90.64%), the first three of which are all numeric identifiers further corroborating our initial findings. Table 1 presents some examples of such over-redaction.

4.4 The MathEd-PII Surrogate Benchmark

By replacing all of the originally redacted PII and newly identified PII with LLM-generated, context-aware surrogates and removing PII annotations for those that were labeled as “Not PII”, we obtained the MathEd-PII surrogate benchmark dataset. A comparison of basic statistics for the source corpus and MathEd-PII is presented in Table 3.

Table 3: PII Statistics Comparison between Source Corpus and MathEd-PII (ordered by the most common PII to the least in MathEd-PII)

Category	Source Corpus	MathEd-PII
Transcripts	1,000	1,000
Messages	115,620	115,620
PII Labels (Total)	5,263	1,995
PERSON	1,915	1,424
URL	245	187
LOCATION	595	121
GRADE_LEVEL	87	107
SCHOOL	88	73
COURSE_NUMBER	1,103	40
NRP	235	25
AGE	0	8
DATE	0	4
US_DRIVER_LICENSE	941	2
PHONE_NUMBER	30	2
IP_ADDRESS	2	2
US_BANK_NUMBER	20	0
US_PASSPORT	2	0

5. PHASE 2: MATH SEGMENTATION AND NUMERIC AMBIGUITY

Since major “Not PII” redactions are numeric identifiers, a natural explanation is *numeric ambiguity*: mathematical content in tutoring transcripts often shares surface-level similarities with numeric PII, making the two difficult to distinguish. Assessing the extent to which numeric ambiguity degrades PII detection—and developing effective mitigation

strategies—therefore requires distinguishing math-dense regions from non-mathematical conversational text.

5.1 Method: Math Density, Semantic Similarity, and Segmentation Optimization

To distinguish math-focused sections in a transcript, we applied a segmentation strategy by first identifying anchor math messages with high density of math vocabulary and then expanding the segment based on semantic similarity [5, 13, 20].

The construction of the math vocabulary set was informed by a comprehensive list of 2,824 K-12 math learning components⁹ curated by the Learning Commons initiative and the Achievement Network [15], which are granular representations of individual math skills and concepts encompassing the Common Core State Standards for Mathematics (CC-SSM) and 18 U.S. state math standards (e.g., California, New York, Texas, and Washington). Some examples of the learning components are “Use conversions to solve multi-step real-world problems”, “Determine the lateral surface area of three-dimensional cylinders in real-world problems”, and “Represent proportional relationships with equations”. The math vocabulary used for math density calculation includes a list of math terms derived from these learning components and regular expressions (RegEx) of common math expressions. The vocabulary can be found here¹⁰.

The **Math Density** (D_{math}) of a message is a weighted sum of single-word math vocabulary terms (V_{math}), multi-word phrases (P_{phrase}), and RegEx patterns ($R_{pattern}$):

$$D_{math}(m) = \frac{1}{|W|} \left[\sum_{w \in W} \mathbb{I}(w \in V) + 1.5 \sum_{p \in P} C(p, m) + 2.0 \sum_{r \in R} C(r, m) \right] \quad (1)$$

where W is the set of tokens in the message, $\mathbb{I}(\cdot)$ is the indicator function, $C(\cdot, m)$ counts Occurrences in message m . The weights were empirically set to attribute higher significance to multi-term concepts and structural math patterns.

A message exceeding a threshold of Math Density, *Anchor Threshold* (T_{anchor}) at 0.05, becomes an anchor message and initializes a math segment. Adjacent messages are added to the math segment if their cosine similarity to the segment centroid (the aggregated semantic meaning of the math context accumulated in the segment computed via `all-MiniLM-L6-v2` embeddings) exceeds a *Similarity Threshold* (T_{sim}) at 0.3. The thresholds were determined via a grid search optimization process around the average math density of the dataset (0.07) from 0.05 to 0.10 in increments of 0.01 and a range of similarity threshold from 0 to 0.5 with a step size of 0.1 with the objective to maximize the capture of “Not PII” (false positives) and minimize the capture of “PII” (true positives) by the math segments. Notably, across all threshold configurations evaluated in the grid search, the proportion of

⁹<https://docs.learningcommons.org/knowledge-graph/understanding-knowledge-graph/about-knowledge-graph>

¹⁰see 2 in https://github.com/ZhuqianZhou/MathEd-PII/blob/69012c0dbde3b3976f36533e2e44f5a5bfdfd718/EDM2026_paper_183_Appendix.pdf

false positives consistently exceeded the proportion of true positives captured by the math segments. This indicates that the finding elaborated below is robust to the choice of threshold. The optimization process and analysis can be seen here¹¹.

5.2 False Positive PII in Math Segments and Numeric Ambiguity

Figure 2 illustrates the distribution of PII categories across mathematical and non-mathematical segments, distinguishing between true positives (validated PII) and false positives (Not PII) arising from numeric ambiguity.

It reveals a stark contrast in detection accuracy between the two domains. In math segments, the data provider’s in-house detection system produced 2,539 false positives compared to only 520 true PII instances—a false-to-true ratio of approximately 4.88:1. In contrast, non-math segments exhibited a much higher precision, with 929 false positives and 1,455 true PII instances (a ratio of 0.64:1).

Specific numeric categories were nearly entirely composed of false positives in mathematical contexts. For instance, as shown in the MATH Segment of Figure 2, categories such as `DATE`, `US_DRIVER_LICENSE`, and `LOCATION` contain substantial numbers of false positives. Interestingly, the `PERSON` category also presents a high number of false positives within math segments. Qualitative review indicates that these names are frequently tied to mathematical word problems (e.g., “If Alex has 5 apples...”) rather than the actual identities of the student or tutor. While over-redacting these fictional entities is perhaps less detrimental to educational utility than over-redacting the mathematical operations themselves, it further illustrates the inability of domain-agnostic systems to distinguish between the *context* of a name (math problem vs. personal disclosure).

This empirical evidence confirms that numeric ambiguity is a structural failure mode in de-identifying math tutoring transcripts: the surface forms of mathematical expressions and problem-related entities systematically overlap with structured personal identifiers, leading generic detection systems to over-redact critical instructional content.

6. PHASE 3: PII DETECTION AND EVALUATION

6.1 PII Detection Methods

6.1.1 Baseline: Microsoft Presidio

As a baseline, we deployed Microsoft Presidio (v2.2) [16]. We utilized its default analyzer which orchestrates a set of predefined recognizers. For high-structure entities (e.g., `EMAIL_ADDRESS`, `URL`, `IP_ADDRESS`, `PHONE_NUMBER`, `US_SSN`, `US_DRIVER_LICENSE`), Presidio employs regular expressions (RegEx). For context-dependent entities (e.g., `PERSON`, `LOCATION`, `NRP`), it leverages a spaCy-based NER model (`en_core_web_lg` and `en_core_web_trf` used here). To address domain-specific requirements, we implemented custom RegEx-based recognizers for `SCHOOL`, `COURSE_NUMBER`, and

¹¹see 3 in https://github.com/ZhuqianZhou/MathEd-PII/blob/69012c0dbde3b3976f36533e2e44f5a5bfdfd718/EDM2026_paper_183_Appendix.pdf

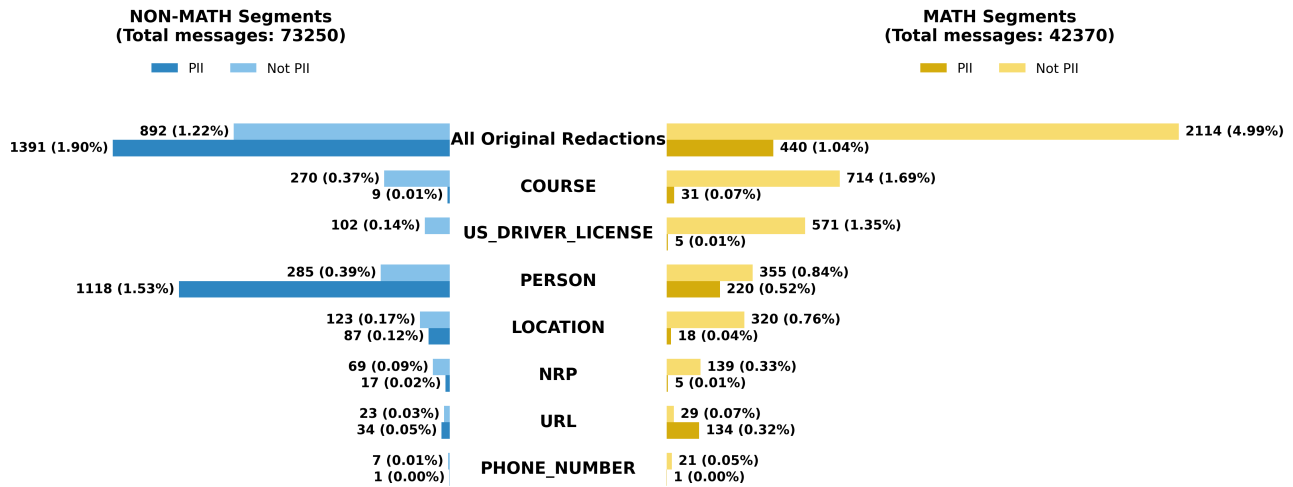


Figure 2: Distribution of Original PII Redactions from the Data Provider’s In-House System (Presidio and Customized Rules) across Different PII Types and Math and Non-Math Segments in terms of true PII and false positives (“Not PII”). Types are ordered by the frequency of “Not PII” in Math Segments. The first seven types are presented.

GRADE_LEVEL. We mapped Presidio’s output entities to the 17 target PII categories.

6.1.2 LLMs using Basic, Math-Aware, and Segment-Aware Prompts

We evaluated the performance of six LLMs—Claude 4.5 Haiku, Claude 4.5 Opus, GPT-4 Turbo, GPT-5.2, Gemini 2.5 Flash, and Gemini 3 Pro—on the task of PII detection to examine the robustness of LLM-based methods. These models were selected because they are widely used and span lightweight, general-purpose, and high-capability model regimes. To ensure data security and compliance, all API requests were routed through Cornell University’s LLM Gateway, LiteLLM¹². This centralized gateway provides a unified interface while enforcing strict privacy protocols; crucially, it operates under a zero-retention policy where payload data is neither logged by the gateway nor used for model training, aligning with institutional IRB requirements.

We first compared model performance under basic prompting against the baseline and selected the higher-performing model from each provider for subsequent evaluation with the math-aware and segment-aware prompting strategies, in order to assess whether these refinements further improve PII detection performance. The three prompt variants are described below.

Basic Prompt: A straightforward instruction acting as a PII specialist. It lists the 17 PII categories and requests the model to identify “ALL PII” in the message, returning a JSON list of text spans and types (exact text here¹³).

Math-Aware Prompt: A refined version of the basic prompt that explicitly addresses the domain challenge. This prompt

aims to reduce false positives stemming from numeric ambiguity by leveraging the model’s general knowledge about math tutoring (exact text here¹⁴).

Segment-Aware Prompt: A refined version of the basic prompt that explicitly points to the math segment label for each message. This prompt aims to reduce false positives stemming from numeric ambiguity through explicit contextual grounding (exact text in here¹⁵).

6.1.3 Evaluation Metrics and Statistical Analysis

We evaluated PII detection performance using standard metrics: Precision, Recall, and the F1-score. Precision measures the proportion of detected text spans that correctly match the ground truth, Recall indicates the proportion of actual PII spans successfully identified, and the F1-score represents their harmonic mean. To account for the inherent variability and potential clustering of errors within different tutoring sessions, we calculated 95% confidence intervals (CIs) using a bootstrap resampling method. We performed 1,000 bootstrap iterations by resampling tutoring transcripts (the primary unit of our corpus) with replacement. For each iteration, metrics were aggregated across the resampled sessions. The 95% CIs reported in our results correspond to the 2.5th and 97.5th percentiles of the resulting bootstrap distribution.

6.2 PII Detection Performance

6.2.1 Overall Performance

Table 4 summarizes the performance of Microsoft Presidio (baseline) and various LLMs using three types of prompts.

The baseline Presidio models achieved high recall (with the

¹²<https://docs.litellm.ai/>

¹³see 4 in https://github.com/ZhuqianZhou/MathEd-PII/blob/69012c0dbde3b3976f36533e2e44f5a5bfdfd718/EDM2026_paper_183_Appendix.pdf

¹⁴see 5 in https://github.com/ZhuqianZhou/MathEd-PII/blob/69012c0dbde3b3976f36533e2e44f5a5bfdfd718/EDM2026_paper_183_Appendix.pdf

¹⁵see 6 in https://github.com/ZhuqianZhou/MathEd-PII/blob/69012c0dbde3b3976f36533e2e44f5a5bfdfd718/EDM2026_paper_183_Appendix.pdf

Table 4: PII Detection Performance Metrics (with 95% Bootstrap Confidence Intervals)

Model	Precision	Recall	F1
<i>Baseline</i>			
Presidio (Large)	0.254 [0.225, 0.285]	0.747 [0.694, 0.794]	0.379 [0.344, 0.415]
Presidio (Transformer)	0.230 [0.208, 0.253]	0.781 [0.731, 0.826]	0.355 [0.328, 0.385]
<i>Basic Prompting</i>			
Claude 4.5 Haiku	0.831 [0.802, 0.860]	0.620 [0.565, 0.675]	0.710 [0.670, 0.752]
Claude 4.5 Opus	0.911 [0.861, 0.946]	0.730 [0.677, 0.783]	0.810 [0.762, 0.853]
Gemini 2.5 Flash	0.848 [0.806, 0.887]	0.771 [0.723, 0.818]	0.807 [0.769, 0.844]
Gemini 3 Pro	0.887 [0.863, 0.910]	0.760 [0.706, 0.812]	0.818 [0.781, 0.851]
GPT-4 Turbo	0.732 [0.695, 0.769]	0.578 [0.530, 0.630]	0.646 [0.609, 0.685]
GPT-5.2	0.759 [0.723, 0.795]	0.707 [0.656, 0.759]	0.732 [0.695, 0.767]
<i>Math-Aware Prompting</i>			
Claude 4.5 Opus	0.905 [0.860, 0.938]	0.720 [0.661, 0.769]	0.802 [0.756, 0.842]
Gemini 3 Pro	0.869 [0.826, 0.906]	0.735 [0.685, 0.787]	0.796 [0.756, 0.835]
GPT-5.2	0.774 [0.740, 0.809]	0.721 [0.670, 0.764]	0.747 [0.711, 0.780]
<i>Segment-Aware Prompting</i>			
Claude 4.5 Opus	0.934 [0.913, 0.951]	0.730 [0.682, 0.784]	0.820 [0.785, 0.855]
Gemini 3 Pro	0.888 [0.862, 0.913]	0.764 [0.713, 0.811]	0.821 [0.788, 0.851]
GPT-5.2	0.800 [0.760, 0.838]	0.720 [0.670, 0.766]	0.758 [0.720, 0.793]

Transformer version reaching a point estimate of 78.1%), but suffered from exceptionally poor precision (spanning 0.23–0.25). This confirms that generic PII detection tools tend to over-redact math-rich content by failing to distinguish mathematical notation from sensitive identifiers. While the point estimates for Presidio’s recall appear high, several top-performing LLMs achieved recall levels that are statistically comparable (with overlapping 95% bootstrap confidence intervals) while providing vastly superior precision and F1 scores. Based on these initial results, Claude 4.5 Opus (Anthropic), Gemini 3 Pro (Google), and GPT-5.2 (OpenAI) were selected as the primary focus for evaluating domain-specific prompt refinements.

Math-aware prompting generally improved the precision of the selected models compared to the basic prompt across most configurations. Segment-aware prompting, which explicitly grounds the detection process in the mathematical context of the message, yielded the highest overall performance for all three models. Notably, Claude 4.5 Opus achieved the highest precision (0.934, 95% CI [0.913, 0.951]) under this condition, while Gemini 3 Pro achieved the highest overall F1 score (0.821, 95% CI [0.788, 0.851]). Crucially, the segment-aware versions of these models maintain recall levels that are on par with the baseline, while reducing false positives by as much as 90%. These results demonstrate that incorporating mathematical discourse context is essential for minimizing over-redaction.

6.2.2 Performance by PII Categories

Given that both the baseline and the top-performing LLMs achieved comparable recall levels (as discussed in

the previous section), we zoom in on category-level performance to investigate how domain-aware prompting reduces false positives to achieve superior precision. Table 5 provides a breakdown of the specific sources of over-redaction, focusing on False Positives (FP) and Precision across the target PII categories under two model configurations: the highest-performing baseline, Presidio LG (using the `en_core_web_lg` SpaCy NER model), and the highest-performing LLM configuration, Gemini 3 Pro with segment-aware prompting.

Table 5: PII Detection False Positives and Precision by Category (ordered by Precision from the highest to the lowest from Gemini 3 Pro with Segment Prompting)

PII Type	Total	Presidio LG		Gemini 3 Pro Seg.	
		FP	Prec	FP	Prec
US_DRIVER_LICENSE	2	344	0.000	0	1.000
PERSON	1422	330	0.775	63	0.949
LOCATION	121	153	0.349	5	0.896
NRP	25	45	0.286	2	0.857
URL	187	21	0.859	21	0.852
SCHOOL	73	17	0.638	10	0.831
GRADE_LEVEL	103	2	0.973	22	0.750
IP_ADDRESS	2	0	1.000	1	0.667
PHONE_NUMBER	2	6	0.250	1	0.667
AGE	8	3	0.000	6	0.538
COURSE_NUMBER	40	46	0.193	43	0.411
DATE	3	3386	0.001	16	0.059
US_BANK_NUMBER	0	5	0.000	-	-

The most striking improvement is observed in numeric categories where the baseline is particularly prone to over-redact mathematical entities. For instance, Presidio generated 3,386 false positives for `DATE` and 344 for `US_DRIVER_LICENSE`, confirming that rule-based systems frequently misinterpret math spans as structured identifiers. In contrast, Gemini 3 Pro with segment-aware prompting reduced `DATE` false positives by over 99.5% (down to 16) and eliminated false positives for `US_DRIVER_LICENSE` entirely. Similarly, for context-heavy categories like `PERSON` and `LOCATION`, the LLM significantly improved precision by reducing the number of false detections. These results underscore that the superior performance of domain-aware LLMs is primarily driven by their ability to selectively filter mathematical noise while maintaining a capture rate on par with standard tools.

6.2.3 Performance by Segment Types

To further investigate how well the LLM addressed numeric ambiguity, we evaluated model performance separately for messages labeled as `NON-MATH` versus `MATH` conversational segments. Table 6 reveals a consistent performance gap across the models, with precision being systematically lower in the `MATH` segment compared to the `NON-MATH` segment.

The stratified analysis in Table 6 provides strong empirical evidence for the numeric ambiguity hypothesis and demonstrates the effectiveness of segment-aware prompting in mitigating it. The baseline Presidio LG model exhibits a big drop in precision when moving from `NON-MATH` (0.346) to

Table 6: PII Detection Performance by Segment Type (NON-MATH vs. MATH)

Model	Prompt	NON-MATH Segments			MATH Segments		
		Prec	Rec	F1	Prec	Rec	F1
Presidio LG	-	0.346	0.756	0.475	0.142	0.720	0.237
Gemini 3 Pro	Basic	0.913	0.819	0.864	0.800	0.591	0.680
Gemini 3 Pro	Segment	0.909	0.815	0.860	0.819	0.618	0.705

MATH segments (0.142), confirming that generic rule-based systems frequently misinterpret mathematical notation as sensitive identifiers. This 59% decline in precision within math-dense regions highlights the risk of over-redaction in tutoring transcripts when domain context is ignored.

In contrast, the Gemini 3 Pro model maintains substantial precision in both contexts, though a performance gap between segments remains. Crucially, the introduction of segment-aware prompting specifically improves the model’s ability to navigate numeric ambiguity within mathematical discourse. In MATH segments, segment-aware prompting increases both precision (0.800 to 0.819) and recall (0.591 to 0.618) compared to basic prompting. This simultaneous gain in both metrics is particularly noteworthy; it suggests that providing the model with an explicit context label allows it to adopt a more nuanced decision boundary—correctly identifying PII that it might otherwise miss (higher recall) while simultaneously avoiding false detections triggered by mathematical symbols (higher precision). The negligible difference in performance within NON-MATH segments further indicates that the benefits of this strategy are structurally concentrated in the regions where numeric ambiguity is most severe.

7. DISCUSSION AND CONCLUSION

This study investigated the challenge of utility-preserving de-identification in the context of math tutoring transcripts, focusing on the phenomenon of numeric ambiguity. By introducing MathEd-PII, the first benchmark dataset for this domain, we provided a rigorous foundation for evaluating PII detection methods that balance privacy protection with instructional fidelity.

Our findings provide strong empirical evidence that numeric ambiguity is a structural failure mode in standard PII detection systems. Generic tools like Microsoft Presidio often misinterpret mathematical expressions as sensitive identifiers, such as dates or driver’s license numbers. Our density-based segmentation analysis revealed that over half of all false positives (55.5%) were concentrated in math-dense regions, which comprise only about a third of the corpus tokens. In these regions, the ratio of false positives to true PII reached nearly 5:1, highlighting a significant risk of removing core instructional content that is essential for educational research.

We demonstrated that LLMs, particularly when guided by domain-aware prompting strategies, significantly outperform traditional baselines. While Presidio achieved a respectable recall, its low precision (F1: 0.380) renders it problematic for creating high-utility datasets. In contrast, Gemini 3 Pro with segment-aware prompting achieved an F1

of 0.821, maintaining high recall while reducing false positives by over 90% in critical numeric categories. The success of segment-aware prompting suggests that providing models with explicit contextual grounding—identifying whether a message is likely to be mathematical—allows them to adopt more nuanced decision boundaries that better distinguish task-content from personal disclosure.

The creation of MathEd-PII through a human-in-the-loop LLM surrogation workflow also offers a methodological contribution. It demonstrates how high-quality benchmark data-sets can be generated even when original, unredacted data cannot be directly shared, by leveraging the generative capabilities of LLMs to create privacy-preserving surrogates that maintain discourse coherence.

However, several limitations and avenues for future work remain. First, while MathEd-PII is the first math-focused PII surrogate benchmark, it is based on 1,000 sessions from a single U.S.-based tutoring platform. As a result, the generalization behavior of the proposed methodology across broader educational domains remains an open question. Expanding this benchmark to include more diverse curricula, grade levels, instructional settings, and international contexts would improve the generalizability of our findings and enable more systematic evaluation across heterogeneous datasets. More broadly, our approach to ensuring applicability across different contexts is to emphasize domain-aware prompting principles that are adaptable to new educational settings, while future work will focus on validating these strategies on additional datasets and subject areas beyond mathematics tutoring.

Second, while this work focused primarily on addressing the precision issues caused by numeric ambiguity, further optimizing recall remains an important avenue for future exploration. Although our domain-aware prompting strategies yielded substantial improvements in F1 and maintained recall levels comparable to standard baselines, a portion of PII still remains undetected. While a detailed investigation into maximizing recall was beyond the scope of this paper, there are several promising directions to bridge this gap, such as: (1) developing multi-stage “cascade” architectures that combine high-recall rule-based systems with high-precision LLM verification; (2) exploring advanced prompting techniques like multi-agent debate to enhance sensitivity to subtle PII mentions; and (3) leveraging the MathEd-PII dataset to fine-tune specialized models that capture domain-specific patterns more effectively and robustly than prompt-only approaches. In addition, the current methodology relies entirely on prompting large language models, which may be brittle across model versions, prompting styles, or unseen distributions. Fine-tuning specialized domain-adapted models therefore represents an important future direction for improving robustness, consistency, and transferability across datasets and deployment settings. Finally, future work could also investigate making these solutions more efficient for low-cost, on-device deployment at educational institutions.

In conclusion, as dialogue-based interaction data becomes increasingly available to educational data mining community, the need for domain-aware de-identification becomes paramount. This work shows that by incorporating mathe-

matical context into the de-identification pipeline, we can move closer to the goal of “utility-preserving privacy”—protecting student and tutor identities without sacrificing the very data that makes educational research possible.

8. ACKNOWLEDGMENTS

This material is based upon initial work completed under National Science Foundation Grant No. 2321499, and support from the Gates Foundation and the Chan Zuckerberg Initiative. Any opinions expressed in this material are those of the authors and do not necessarily reflect funders’ views.

9. DATA AVAILABILITY

The MathEd-PII dataset is publicly available at <https://huggingface.co/datasets/NationalTutoringObservatory/MathEd-PII>.

10. REFERENCES

- [1] A. Caines, H. Yannakoudakis, H. Allen, P. Pérez-Paredes, B. Byrne, and P. Buttery. The teacher-student chatroom corpus version 2: more lessons, new annotation, automatic detection of sequence shifts. In D. Alfter, E. Volodina, T. François, P. Desmet, F. Cornillie, A. Jönsson, and E. Rennes, editors, *Proceedings of the 11th Workshop on NLP for Computer Assisted Language Learning*, pages 23–35, Louvain-la-Neuve, Belgium, Dec. 2022. LiU Electronic Press.
- [2] D. S. Carrell, B. Malin, J. Aberdeen, S. Bayer, and C. Clark. Hiding in plain sight: Use of realistic surrogates to reduce exposure of protected health information in clinical text. *Journal of the American Medical Informatics Association*, pages 342–348, 2013.
- [3] G. Deacon and G. Chojnacki. Impacts of upchieve on-demand tutoring on students’ math knowledge and perceptions. middle years math grantee report series. *Mathematica*, 2023.
- [4] Federal Trade Commission. Children’s online privacy protection act (coppa) guidance. <https://www.ftc.gov/business-guidance/resources/complying-coppa-frequently-asked-questions>, 2013. Accessed: 2026-02-10.
- [5] M. Galley, K. McKeown, J. Hirschberg, and E. Shriberg. Discourse segmentation of multi-party conversation. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 562–569. Association for Computational Linguistics, 2003.
- [6] S. L. Garfinkel. De-identification of personal information. NIST Interagency Report 8053, National Institute of Standards and Technology, Oct. 2015.
- [7] S. L. Garfinkel. De-identifying government data sets: Techniques and governance. NIST Special Publication 800-188, National Institute of Standards and Technology, Sept. 2023.
- [8] O. Henkel, H. Horne-Robinson, N. Kozhakhmetova, and A. Lee. Effective and scalable math support: Evidence on the impact of an ai-tutor on math achievement in ghana. *arXiv preprint arXiv:2402.09809*, 2024.
- [9] L. Holmes, S. Crossley, N. Hayes, D. Kuehl, A. Trumbore, and G. Gutu-Robu. De-identification of student writing in technologically mediated educational settings. In *Polyphonic Construction of Smart Learning Ecosystems: Improving Inclusive Digital Education*, pages 177–189, Singapore, 2023. Springer Nature Singapore.
- [10] L. Holmes, S. Crossley, H. Sikka, and W. Morris. Piilo: an open-source system for personally identifiable information labeling and obfuscation. *Information and Learning Sciences*, 124(9-10):266–284, Oct. 2023.
- [11] L. Holmes, S. Crossley, J. Wang, and W. Zhang. The cleaned repository of annotated personally identifiable information. In P. Benjamin and D. E. Carrie, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 790–796, Atlanta, Georgia, USA, July 2024.
- [12] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd. spacy: Industrial-strength natural language processing in python. <https://doi.org/10.5281/zenodo.1212303>, 2020.
- [13] Z. Huang, W. Xu, and K. Yu. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv: 1508.01991*, 2015.
- [14] K. R. Koedinger, R. S. Baker, K. Cunningham, A. Skogsholm, B. Leber, and J. Stamper. A data repository for the edm community: The pslc datashop. *Handbook of educational data mining*, 43:43–56, 2010.
- [15] Learning Commons Initiative. Learning components. <https://docs.learningcommons.org/knowledge-graph/entity-and-relationship-reference/learnin-g-components>, 2024. Accessed: 2026-02-05.
- [16] Microsoft. Presidio: Data protection and de-identification sdk. <https://github.com/microsoft/presidio>, 2020. GitHub repository. Accessed 2026-01-20.
- [17] Microsoft. presidio-research: Research utilities for Presidio (including synthetic text generation). <https://github.com/microsoft/presidio-research>, 2021. GitHub repository. Accessed 2026-01-20.
- [18] M. C. Mihaescu and P. S. Popescu. Review on publicly available datasets for educational data mining. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 11(3):e1403, 2021.
- [19] K. Mishra, H. Pagare, and K. Sharma. A hybrid rule-based NLP and machine learning approach for PII detection and anonymization in financial documents. *Scientific Reports*, 15(1):22729, 2025.
- [20] H. Misra, F. Yvon, O. Cappé, and C. Jose. Text segmentation: A topic modeling perspective. *Information Processing & Management*, 47(4):528–544, 2011.
- [21] I. Neamatullah, M. M. Douglass, L.-W. H. Lehman, A. Reisner, M. Villarroel, W. J. Long, P. Szolovits, G. B. Moody, R. G. Mark, and G. D. Clifford. Automated de-identification of free-text medical records. *BMC Medical Informatics and Decision Making*, 8(32), 2008.
- [22] E. Prihar, A. Moore, and N. Heffernan. Identifying explanations within student-tutor chat logs. In *Proceedings of the 15th International Conference on Educational Data Mining*, page 773, 2022.
- [23] M. Savkin, T. Ionov, and V. Kononov. SPY: Enhancing privacy with synthetic PII detection

- dataset. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*, pages 236–246. Association for Computational Linguistics, 2025.
- [24] Y. Shen, Z. Ji, J. Lin, and K. R. Koedinger. Enhancing the de-identification of personally identifiable information in educational data. *arXiv preprint arXiv: 2501.09765*, 2025.
- [25] S. Singhal, A. F. Zambrano, M. Pankiewicz, X. Liu, C. Porter, and R. S. Baker. De-identifying student personally identifying information with gpt-4. In *Proceedings of the 17th International Conference on Educational Data Mining (EDM)*, pages 559–565, 2024.
- [26] N. Slijepcevic and A. Yaylali. Leveraging “khanmigo” generative ai-powered tool for personalized tutoring to learn scientific concepts. *Journal of Teaching and Learning*, 19(4):155–178, 2025.
- [27] A. Stubbs and Ö. Uzuner. Automated systems for the de-identification of longitudinal clinical narratives: Overview of the 2014 i2b2/uthealth shared task track 1. *Journal of Biomedical Informatics*, 58:S11–S19, 2015.
- [28] R. Tang, X. Han, X. Jiang, and X. Hu. Does synthetic data generation of LLMs help clinical text mining? *arXiv preprint arXiv: 2303.04360*, 2023.
- [29] U.S. Department of Education, Privacy Technical Assistance Center (PTAC). Data de-identification: An overview of basic terms. https://studentprivacy.ed.gov/sites/default/files/resource_document/file/data_deidentification_terms_0.pdf, 2012. Updated May 2013. Accessed 2026-01-20.
- [30] U.S. Department of Health and Human Services, Office for Civil Rights. Guidance regarding methods for de-identification of protected health information in accordance with the HIPAA privacy rule. https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/understanding/coveridentities/De-identification/hhs_deid_guidance.pdf, 2012. Accessed 2026-01-20.
- [31] J. Zambrano, K. Singhal, L. Pankiewicz, R. Baker, L. Porter, and L. Liu. De-identifying student personally identifying information in discussion forum posts with large language models. *Information and Learning Sciences*, 126(5/6):401–424, 2025.
- [32] M. Zent, D. Smith, and S. Woodhead. PIIvot: A lightweight NLP anonymization framework for question-anchored tutoring dialogues. *arXiv*, 2025.