

# Analysing Knowledge Construction in Online Learning: Adapting the Interaction Analysis Model for Unstructured Large-Scale Discourse

Jindi Wang<sup>a</sup>, Yidi Zhang<sup>b</sup>, Zhaoxing Li<sup>\*c</sup>, Pedro Bem-Haja<sup>b</sup>, Ioannis Ivrissimtzis<sup>a</sup>, Zichen Zhao<sup>d</sup> and Sebastian Stein<sup>c</sup>

<sup>a</sup>Department of Computer Science, Durham University, Durham, UK

<sup>b</sup>Department of Education and Psychology, University of Aveiro, Aveiro, Portugal

<sup>c</sup>School of Electronics and Computer Science, University of Southampton, Southampton, UK

<sup>d</sup>School of Mathematics and Statistics, Lanzhou University, Lanzhou, China

---

## ARTICLE INFO

### Keywords:

Knowledge Construction  
Educational Discourse Analysis  
Automated Text Classification  
Transformer-based Language Models

## Abstract

The rapid expansion of online courses and social media has generated large volumes of unstructured learner-generated text. Understanding how learners construct knowledge in these spaces is crucial for analysing learning processes, informing content design, and providing feedback at scale. However, existing approaches typically rely on manual coding of well-structured discussion forums, which does not scale to the fragmented discourse found in online learning. This study proposes and validates a framework that combines a codebook inspired by the Interaction Analysis Model with an automated classifier to enable large-scale analysis of knowledge construction in unstructured online discourse. We adapt four comment-level categories of knowledge construction: Non-Knowledge Construction, Share, Explore, and Integrate. Three trained annotators coded a balanced sample of 20,000 comments from YouTube education channels. The codebook demonstrated strong reliability, with Cohen's  $\kappa = 0.79$  on the main dataset and 0.85–0.93 across four additional educational domains. For automated classification, bag-of-words baselines were compared with transformer-based language models using 10-fold cross-validation. A DeBERTa-v3-large model achieved the highest macro-averaged F1 score (0.841), outperforming all baselines and other transformer models. External validation on four domains yielded macro-F1 above 0.705, with stronger transfer in medicine and programming, where discourse was more structured and task-focused, and weaker transfer in language and music, where comments were more varied and context-dependent. Overall, the study shows that theory-driven, semi-automated analysis of knowledge construction at scale is feasible, enabling the integration of knowledge-construction indicators into learning analytics and the design of online learning environments.

---

## 1. Introduction

Identifying and supporting knowledge construction in online discussions remains a significant challenge in education (Gunawardena et al., 2023). As online learning environments such as massive open online courses, discussion forums, and social media continue to expand, learners are producing large amounts of unstructured textual data (Dubovi & Tabak, 2020; Nguyen & Diederich, 2023). Understanding how knowledge is constructed within these environments provides an important foundation for analysing the cognitive processes of learners (Gunawardena et al., 1997; Henri, 1992). For example, such insights can inform the design of more targeted instructional materials and scaffolds that promote higher-level epistemic moves (Scardamalia & Bereiter, 2006; Weinberger & Fischer, 2006), support learning analytics indicators that monitor the quality of discourse over time rather than only counting participation (Gašević et al., 2015; Wise et al., 2013), and help instructors provide more timely and focused feedback by identifying when learners are merely sharing information versus exploring or integrating ideas (Dubovi & Tabak, 2020; Wever et al., 2006).

Traditional analyses of knowledge construction have relied on manual content analysis. For example, Henri

(1992) proposed a five-dimensional framework for computer-mediated communication that includes participation, interaction, social, cognitive, and metacognitive dimensions, providing a systematic basis for interpreting learning processes in message transcripts. Gunawardena et al. (1997) introduced the Interaction Analysis Model (IAM), which conceptualises the process of knowledge construction in computer-mediated communication as a progressive sequence of phases, from sharing and comparing information to negotiating meaning, constructing new knowledge, and reaching mutual understanding. The IAM has since become a foundational framework for examining collaborative learning and knowledge building in online environments (Lucas et al., 2014). Subsequent studies further refined the analysis of knowledge construction. Hmelo-Silver (2003) employed mixed methods to investigate how groups build shared understanding, demonstrating that detailed conversational exchanges and broader thematic patterns characterise computer-supported collaborative learning. Weinberger and Fischer (2006) proposed a multidimensional framework for argumentative knowledge construction, distinguishing participation, epistemic, argumentative, and social dimensions to capture how learners co-construct knowledge through discourse. Together, these studies have advanced methodological approaches for examining knowledge construction in digital learning contexts.

---

ORCID(s):

Despite recent progress, manual coding remains labour-intensive, time-consuming, and susceptible to subjectivity (Dornauer et al., 2024). Studies must establish and maintain inter-rater reliability, and disagreements often require negotiated coding (Hallgren, 2012). To reduce delay and improve scalability, researchers have explored automated discourse classification for timely feedback. Deep learning and other supervised approaches have been used to identify phases of knowledge construction (Hu et al., 2021). However, performance varies with language, context, and the availability of labelled data, and key challenges persist (Hu et al., 2022).

Recent advances in large pre-trained language models have substantially improved contextual representation in educational discourse analysis, including work grounded in the Interaction Analysis Model (IAM) (Castellanos-Reyes et al., 2025; Gunawardena et al., 2023; Sakeef et al., 2025). Transformer architectures such as BERT, RoBERTa, and DeBERTa have demonstrated strong performance on related discourse tasks (e.g., stance detection, sentiment analysis, and argumentation modelling) (Devlin et al., 2019; He, Liu, et al., 2021). However, their application to fine-grained classification of IAM-based knowledge construction remains limited. Prior studies typically rely on small or imbalanced datasets, employ coarse category schemes, and focus on single, well-structured discussion contexts, with little evaluation of cross-domain robustness (Dubovi & Tabak, 2020; Gunawardena et al., 2023; Nguyen & Diederich, 2023).

In addition, models optimised for specific datasets or domains are often quickly outperformed by newer language-model variants, improving raw performance but hindering the development of stable, theory-grounded approaches that can be reused across studies (Castellanos-Reyes et al., 2025; Mim et al., 2025). Consequently, existing automated methods provide only partial support for reliably distinguishing nuanced IAM phases—such as sharing, exploration, and integration—within large-scale, noisy social media discourse (Dubovi & Tabak, 2020; Nguyen & Diederich, 2023).

Therefore, this study aims to propose and validate a framework that integrates an IAM-inspired codebook with an automated classifier trained on a balanced category training set, enabling large-scale analysis of knowledge construction in online learning environments.

The remainder of this section provides the conceptual background for the study. Section 1.1 discusses how knowledge construction manifests in online learning discourse, Section 1.2 reviews the Interaction Analysis Model and our rationale for adapting it to comment-level analysis, and Section 1.3 surveys work on automated classification of knowledge construction and related constructs. Section 1.4 then states the overall aim of the study and formulates the research questions. Section 2 describes the data, annotation procedure, and modelling pipeline; Section 3 presents the experimental results; Section 4 discusses implications and limitations; and Section 5 concludes.

## 1.1. Identifying knowledge construction in online learning discourse

Knowledge construction is a cognitive process in which learners actively interpret, organise, and integrate information to build new understanding (Scardamalia & Bereiter, 2006). In online learning environments, this process becomes visible through learners' discourse and comment interactions, where they articulate claims, question ideas, provide justifications, and integrate multiple perspectives (Mayer, 1996; Nguyen & Diederich, 2023). Such forms of engagement further encourage learners to set their own goals and regulate their learning (Greenhow & Lewin, 2016; Tazhenova et al., 2024).

However, discourse and comment interaction in computer-mediated communication is not always linear. Messages may appear as independent contributions rather than as parts of a structured exchange (Maccoccia, 2004). In many online environments, discussions are asynchronous and lack reciprocal engagement, making it difficult to maintain coherent conversational threads (Krykoniuk et al., 2025). Participants may misinterpret the conversational structure and it is often unclear who is responding to whom (Maccoccia, 2004). This pattern is even more evident on social media platforms, such as YouTube, where most comments are short, fragmented, and do not form extended exchanges (Tanskanen, 2021).

Despite this structural fragmentation, individual comments often still contain linguistically identifiable elements that correspond to the different phases of knowledge construction such as triggering events, exploration, integration, and resolution, even though these phases do not appear in a strict sequential order (Garrison et al., 2001). Research in epistemic cognition and argumentation further demonstrates that such discursive moves can reliably indicate forms of knowledge construction even when direct reciprocal interaction is absent (Dubovi & Tabak, 2020; Nguyen & Diederich, 2023). Thus, although asynchronous comment streams lack the structural coherence needed for phase-by-phase interactional analysis, they nevertheless provide analytically meaningful traces of knowledge construction processes (Ba et al., 2023; Gunawardena et al., 2023).

Moreover, although individual comments often appear structurally independent, they may still respond to video content or to the contributions of other viewers (Krykoniuk et al., 2025). Even without explicit reply chains, comment sections can display coherence and cross-modal cohesion (Tanskanen, 2021). From this perspective, online comment spaces can be understood as multi-threaded dialogic networks in which individual contributions are part of a broader shared discourse and a collective cognitive space (Kimmerle et al., 2015; Ye & Pennisi, 2022).

Therefore, although knowledge construction can be identified within individual comments, the loose and fragmented structure of social media content often makes such identification inconsistent. This indicates the need for more systematic and scalable approaches to improve the reliability of identifying knowledge construction in online learning contexts.

## 1.2. Using the IAM to analyse knowledge construction

The IAM is widely used to examine knowledge construction in online learning (Gunawardena et al., 1997). It delineates five phases through which learners progressively build knowledge: 1) sharing and comparing information, 2) discovering and exploring dissonance or inconsistency among ideas, concepts or statements, 3) negotiating meaning/co-constructing knowledge, 4) testing and modifying a proposed synthesis or co-construction, and 5) agreeing with statement(s)/applying newly constructed meaning. Although the IAM was originally conceptualised as a structured and progressive interactional process (Lucas et al., 2014), the IAM's phases can also be interpreted as epistemic functions that may surface within individual contributions (Gunawardena et al., 2023). This reinterpretation is particularly relevant in online environments where discourse is often non-linear and fragmented (Krykoniuk et al., 2025).

Applying the IAM to single comments therefore entails a conceptual shift. Comment-level coding captures the epistemic functions enacted in an individual contribution rather than the interactional progression theorised by Gunawardena et al. (1997). This distinction is central to construct validity, while the adapted categories provide evidence of cognitive engagement, they no longer represent the temporal or collaborative trajectory of co-construction (Nguyen & Diederich, 2023). This interpretation aligns with a broader methodological shift in online discourse analysis, where epistemic moves are examined as linguistic indicators that need not follow a linear sequence (Gunawardena et al., 2023). From this perspective, even fragmented or weakly threaded digital environments contain analytically meaningful traces of cognitive work, despite the absence of observable interactional progression (Dubovi & Tabak, 2020).

This approach affords both strengths and limitations. Conceptually, treating discourse as evidence of how learners publicly display their thinking connects comment-level analysis to interaction-oriented models (Garrison et al., 2001; Hmelo-Silver, 2003). For example, Dubovi and Tabak (2020) observed that comments on YouTube science videos often progressed from simple information exchange to argumentative negotiation, which suggests the potential for deeper knowledge construction. In a similar way, Nguyen and Diederich (2023) adapted the IAM to remove its dependence on linear reply structures and found that in TikTok educational content, comments were predominantly social, with knowledge construction emerging mainly through opinion sharing and the exploration of disagreement.

At the same time, reducing the unit of analysis to single comments inevitably constrains the granularity of the model (Gunawardena et al., 1997). Higher-level processes such as negotiating meaning or co-construction are harder to infer without interactional context (Lucas et al., 2014), and empirical work shows that these phases occur infrequently even in structured online discussions (Gunawardena et al., 2023). Moreover, increases in interaction frequency do not necessarily lead to higher levels of knowledge-construction

phases (Hew & Cheung, 2011). Consequently, work in non-interactive environments has often collapsed later IAM phases into broader, lower-level categories to improve coding reliability (Nguyen & Diederich, 2023). This pragmatic adjustment does not diminish the value of the IAM; rather, it situates comment-level analyses within an emerging paradigm in which IAM phases function as discourse markers rather than as strict interactional stages (Dubovi & Tabak, 2020).

Therefore, although the IAM can be effectively applied at the comment level to detect traces of knowledge construction, such use requires careful reconceptualisation of the phases. Ideally, it should also be supported by complementary analytical techniques that strengthen the interpretive validity of phase assignments.

## 1.3. Automated classification of knowledge construction

Recent advances in natural language processing have enabled the automated classification of knowledge construction in online learning discourse (Gunawardena et al., 2023). Early work demonstrated that neural models can identify discourse segments associated with different phases of cognitive presence, establishing the feasibility of such approaches (Hu et al., 2021). More recently, transformer-based architectures such as BERT (Devlin et al., 2019) and DeBERTa (He, Liu, et al., 2021) have achieved substantial improvements in fine-grained discourse classification, showing strong agreement with human annotations in educational contexts (Ba et al., 2023). Extensions incorporating multi-label classification have further improved robustness by capturing the overlapping and fluid nature of epistemic phases (Hu et al., 2023). Studies integrating behavioural or semantic metadata indicate, however, that textual features remain the primary predictive signal (Dornauer et al., 2024). At the same time, recent reviews highlight the need to balance performance gains from large language models with interpretability and cross-context robustness (Castellanos-Reyes et al., 2025).

Despite these advances, automated IAM-based classification faces persistent structural and methodological limitations. The nonlinear and weakly threaded nature of many online platforms—such as YouTube, Facebook, and Twitter/X—obscures conversational context and constrains interaction-level analysis (Aragón et al., 2017; Dubovi & Tabak, 2020; Herring, 1999; Naab et al., 2023). Beyond structural issues, construct validity remains a central concern: machine-learning models infer IAM categories from statistical regularities in text, even though these categories represent cognitive functions that are not directly observable. This creates a risk of reliance on superficial lexical cues, particularly when datasets are small, imbalanced, or domain specific (Breivik, 2016; Dornauer et al., 2024; Hu et al., 2021; Kovanović et al., 2016; Neto et al., 2021). Ensuring validity therefore requires diverse training data, safeguards against overfitting, and evaluation strategies that test alignment with theoretical expectations rather than surface linguistic patterns (Ba et al., 2023; Ferreira et al., 2020).

To accommodate fragmented discourse structures, recent studies have applied IAM coding at the level of individual comments. Gunawardena et al. (2023) reported modest performance on a small corpus of discussion posts, while Nguyen and Diederich (2023) achieved higher F1 scores using BERT on TikTok comments but relied on uneven label distributions and a single-domain setting. Overall, existing work remains limited in scale and scope, with little evidence for cross-domain generalisation. Consequently, larger and more balanced datasets, combined with transformer-based models and external validation across educational domains, are needed to establish reliable and reusable approaches for automated knowledge-construction analysis.

#### 1.4. Study aim and contributions

Against this background, this study operationalises the Interaction Analysis Model (IAM) for large-scale, non-structured online comment streams, where contributions are typically short, non-sequential, and weakly threaded (e.g., YouTube comments). We combine a manually developed, IAM-inspired codebook with supervised classification: a balanced, human-annotated sample provides ground truth for training and evaluating automated classifiers that can then be applied at scale to unlabelled comments.

We make three contributions. First, we adapt IAM to a compact four-category, comment-level scheme (*NonKC*, *Share*, *Explore*, *Integrate*) tailored to non-interactive, unstructured discourse and demonstrate its reliability across domains. Second, we construct and (on acceptance) release a 20,000-comment, manually annotated corpus and use it to benchmark classical and transformer-based classifiers, showing the practical trade-offs between accuracy and stability. Third, we evaluate cross-domain generalisation by validating the best model on four additional educational domains (language, medicine, music, programming) and analysing per-category transfer behaviour.

These aims are operationalised in three research questions:

- RQ1: How can the IAM be adapted into a reliable comment-level coding scheme for non-structured, nonlinear online comments?
- RQ2: On a large, manually annotated corpus, how do traditional machine-learning and transformer-based models compare in identifying IAM categories, and which configuration balances accuracy and stability?
- RQ3: To what extent does the best-performing classifier (DeBERTa-KC) generalise to other educational domains?

## 2. Methods

Figure 1 presents an overview of our end-to-end workflow, spanning data collection through model evaluation and reproducibility. The pipeline is organised into five stages: *Collect*, *Capture*, *Ingest*, *Compute*, and *Store & Use*. It is implemented through modular Python scripts with integrated experiment tracking, and the full source code is available in

an anonymised GitHub repository. The repository link will be released upon acceptance to preserve double-blind review. All key decisions in this pipeline (tokenisation, fold splits, hyperparameters, and metrics) are systematically versioned to ensure auditability and to enable exact reruns of every experiment, thereby strengthening reproducibility.

#### 2.1. Data collection and preprocessing

We collected publicly available YouTube comments using the YouTube Data API v3, in accordance with the platform's Terms of Service, with ethical approval from the University of xxx (institutional details withheld for double-blind review). Comments were retrieved from four science education channels—Museum of Science, SciShow, Sick Science!, and World Science Festival—spanning 2022–2024 and including both *Videos* and *Shorts*. Channels were selected for their broad content and large subscriber bases (>750,000) with extensive video libraries (>900 uploads), ensuring a diverse comment pool.

The dataset reflects moderately structured educational discourse characterised by relatively high content quality and sustained audience engagement, rather than the full spectrum of YouTube interactional styles. We collected 2,198 videos, yielding 609,763 top-level comments and 154,180 replies. All collection scripts, channel identifiers, and cut-off dates will be provided in an anonymised GitHub repository upon acceptance.

For external validation, we sampled four additional learning channels in Language (BBC Learning English), Medicine (Intellect Medicos), Music (Learn Piano with Jazer Lee), and Programming (Programming with Mosh), totalling 42,152 comments. These domains exhibit distinct epistemic and communicative profiles, providing a robust test of model generalisation.

Preprocessing involved removing duplicate comments, lowercasing, compressing whitespace, and retaining discourse-relevant punctuation. Tokenisation used SentencePiece to ensure compatibility with the DeBERTa backbone, and sequences were truncated or padded to 256 tokens.

To prevent content leakage, cross-validation folds were assigned at the video level, so all comments from a given video were used exclusively for training or validation. Ten folds were stratified by video format (*Shorts* vs. regular *Videos*) to ensure balanced evaluation and reliable estimates of model performance.

#### 2.2. Knowledge construction coding

Prior to annotation, three coders were instructed using the codebook (see Table 1 and 2). They then independently annotated a random sample of 1,000 platform comments presented in random order, without any discussion among them. The resulting Cohen's  $\kappa$  was 0.67. Coders were encouraged to provide brief written notes for borderline cases. Following a discussion of discrepancies, boundary conditions were incorporated into the codebook (see Table 10 and Table 11 in the appendix). The coders then independently annotated a second random sample of 1,000 comments, drawn

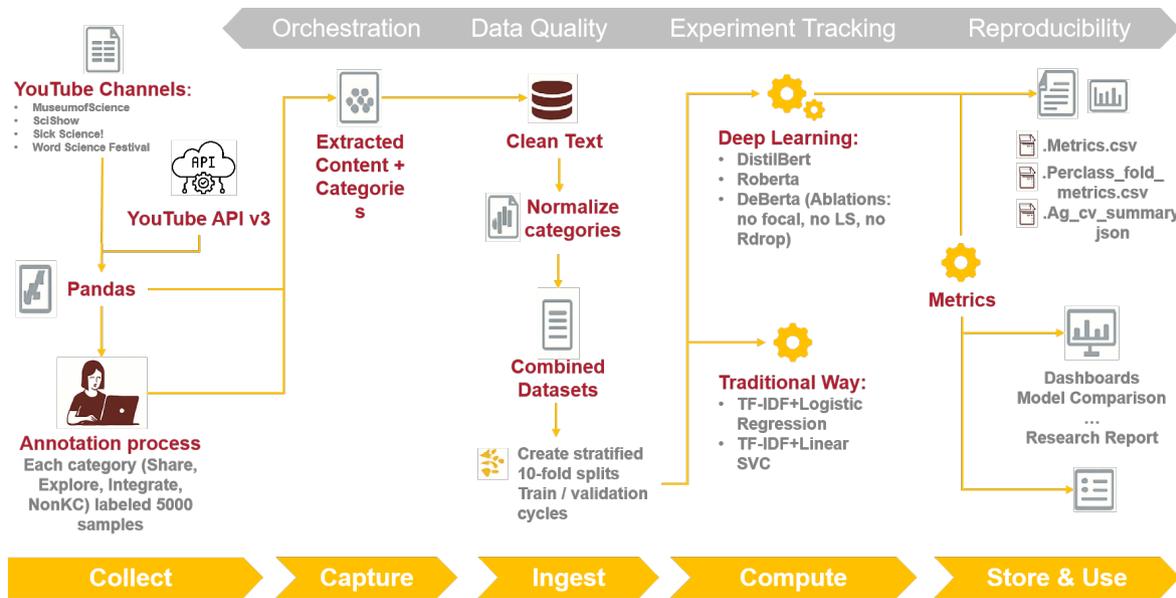


Figure 1: Data collection and model evaluation pipeline.

from a different set than the first round, after which the Cohen’s  $\kappa$  increased to 0.88.

To construct the main annotated dataset, we first addressed the substantial class imbalance in the initial sample. Because *NonKC* and *Share* occurred far more frequently than *Explore* and *Integrate*, coder A continued annotating until the minority categories reached 6,000 instances, resulting in a preliminarily oversampled dataset. Subsequently, comments with duplicate content were removed, and a subset of 5,700 comments per category was created. All labels assigned by coder A were then removed, and coders B and C independently re-annotated these comments. At this stage, inter-rater reliability improved to a Cohen’s  $\kappa$  of 0.79, indicating a high level of agreement. When coders B and C reached agreement, their label was retained as final; when disagreement occurred, the label matching coder A’s earlier annotation was adopted. Remaining disagreements were resolved through discussion. Finally, a balanced random sample of 5,000 comments per category (20,000 comments in total) was drawn to form the main analysis set.

After completing the main annotated dataset, we applied the same codebook to create test sets for external cross-domain validation. During this process, coders B and C were informed of the learning domain of each video to prevent domain-specific content features from being misclassified. Each domain contributed two subsets. First, 600 comments were randomly sampled and double-annotated to form a random test set reflecting the natural distribution of knowledge-construction categories. Second, we targeted at least 100 instances of each category in every domain by oversampling underrepresented categories from the remaining comments; when a category contained fewer than 100 comments, all available instances were included. Classifier predictions were used to identify likely candidates, which were then annotated

manually. This procedure yielded an augmented subset with sufficient coverage for per-category analysis.

While oversampling ensures adequate per-category coverage, it necessarily alters the natural distribution of knowledge-construction categories within each domain. As such, these augmented sets should be interpreted as diagnostic test beds rather than representative domain corpora. This design choice strengthens per-class evaluation but reduces comparability with real-world class frequencies, a trade-off common in educational NLP when minority epistemic categories are infrequent (Ba et al., 2023; Hu et al., 2021; Neto et al., 2021).

As noted by Gunawardena et al. (2023), a single comment may reflect more than one developmental stage. However, the difference between assigning multiple stages and assigning only the highest stage is not statistically significant. Based on this finding, Gunawardena et al. (2023) argued that multilabel annotation is unnecessary when sufficient training data are available. Accordingly, in this study, each comment in the annotated dataset was assigned the label corresponding to the highest applicable category.

### 2.3. Models and baselines

We benchmark **DeBERTa-KC** against two classical and three transformer baselines: TF-IDF (character  $n$ -grams 3–5) + Logistic Regression (balanced class weights) and TF-IDF + LinearSVC, both implemented with scikit-learn (Pedregosa et al., 2011); DistilBERT-base (Sanh et al., 2019), RoBERTa-base (Liu et al., 2019), and DeBERTa-v3-base (He, Gao, & Chen, 2021). All transformer models are implemented using the Hugging Face Transformers library (Wolf et al., 2020), and all models share identical tokenisation, fold assignments, and optimisation protocol for comparability. The proposed **DeBERTa-KC** architecture builds upon the *DeBERTa-v3-large* backbone. Figure 2 illustrates the overall architecture.

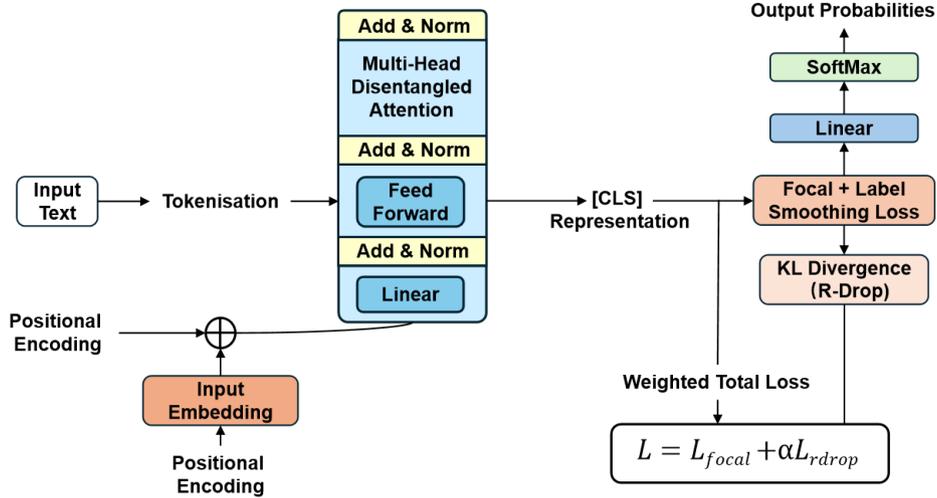


Figure 2: Proposed DeBERTa-KC model.

### 2.3.1. Backbone encoder and classification head

We employ **microsoft/deberta-v3-large**, a transformer with 24 layers, 16 self-attention heads per layer, and hidden size 1024. DeBERTa disentangles content and positional embeddings and applies a relative position bias in self-attention, improving contextual representation over BERT and RoBERTa for fine-grained discourse signals (He, Liu, et al., 2021).

Each input comment is tokenised (SentencePiece), truncated or padded to a maximum of 256 tokens, and fed to a DeBERTa-v3-large encoder with the standard dropout configuration (dropout probability  $p = 0.1$  in the encoder and classification head during training). We use the final-layer [CLS] representation  $h_{\text{CLS}}$  as a sequence summary and, during training, apply this dropout before the linear classification head to reduce overfitting; dropout is disabled at inference time:

$$\mathbf{z} = \text{Dropout}(h_{\text{CLS}}), \quad \mathbf{p} = \text{Softmax}(\mathbf{W}\mathbf{z} + \mathbf{b}),$$

where  $\mathbf{W} \in \mathbb{R}^{4 \times 1024}$  and  $\mathbf{b} \in \mathbb{R}^4$ . The vector  $\mathbf{p}$  yields class probabilities for the four KC categories.

A compact head (single linear layer) minimises additional parameters, reducing overfitting on rare but deeper knowledge-construction labels (e.g., *Integrate*) while preserving transparency, as the head's logits can be inspected or temperature-scaled for calibrated advisory use in learning analytics tools.

### 2.3.2. Training objective

To improve robustness under imbalance and promote trustworthy probabilities, the loss integrates three complementary components: *Focal Loss*, *Label Smoothing*, and *R-Drop*. Together they address (a) class-frequency skew, (b) overconfident predictions, and (c) stochastic inconsistency.

We extend cross-entropy with Focal weighting (Lin et al., 2020) and Label Smoothing (Szegedy et al., 2016). For gold

label  $y$  and predicted probabilities  $\mathbf{p}$ :

$$\mathcal{L}_{\text{FL}} = -(1 - p_y)^\gamma \log(p'_y), \quad p'_y = (1 - \epsilon)p_y + \frac{\epsilon}{K},$$

with focusing parameter  $\gamma = 2.0$ , smoothing  $\epsilon = 0.05$ , and  $K = 4$  classes. Focal Loss down-weights easy examples and gives more gradient signal to hard cases (often *Explore* or *Integrate*), while smoothing curbs overconfidence for better-calibrated probabilities.

R-Drop encourages consistency between two stochastic forward passes of the same input by minimising bidirectional KL divergence (Liang et al., 2021). For two predicted distributions  $\mathbf{p}_1$  and  $\mathbf{p}_2$ :

$$\mathcal{L}_{\text{RD}} = \frac{1}{2} \left[ \text{KL}(\mathbf{p}_1 \parallel \mathbf{p}_2) + \text{KL}(\mathbf{p}_2 \parallel \mathbf{p}_1) \right].$$

This reduces variance due to dropout and attention noise, which is important when model outputs underpin educator workflows (e.g., alerting on threads needing facilitation).

The final loss is an unweighted sum of the focal-loss and R-Drop terms:

$$\mathcal{L} = \mathcal{L}_{\text{FL}} + \lambda_{\text{RD}} \mathcal{L}_{\text{RD}}, \quad \lambda_{\text{RD}} = 1.0.$$

We fix  $\lambda_{\text{RD}}$  to 1.0, rather than tuning it, to keep the configuration simple, and backpropagate gradients through the encoder and classification head jointly.

### 2.3.3. Training procedure

We fine-tune with AdamW and cosine decay (initial learning rate  $1 \times 10^{-5}$ , weight decay 0.05, warm-up ratio 0.1), early stopping (patience = 2), and mixed precision (FP16/BF16). Batch sizes are 8 (train) and 16 (evaluation). Stratified 10-fold cross-validation ensures balanced label representation.

At inference, we form an ensemble by averaging the predicted class probabilities across folds, using simple (unweighted) averaging so that all fold models are treated symmetrically, and the combination method remains transparent and easily comparable across experiments:

$$\hat{\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N \mathbf{p}^{(i)}, \quad N = 10.$$

The predicted label is  $\arg \max_c \hat{p}_c$ . Ensembling reduces variance and stabilises per-class metrics across folds. Alongside  $\hat{\mathbf{p}}$ , we compute predictive entropy  $H(\hat{\mathbf{p}}) = -\sum_c \hat{p}_c \log \hat{p}_c$  as a lightweight uncertainty score.

We implement the model in PyTorch via HuggingFace Transformers, enabling reproducible configuration of tokenisation, truncation length (256), optimiser, and schedulers. All random seeds are fixed; training and evaluation artefacts (metrics per fold, confusion matrices, hyperparameters) are versioned.

#### 2.3.4. Evaluation protocol and metrics

We use *stratified, grouped-by-video* 10-fold cross-validation on the 20,000-item development corpus: in each run, we train on nine folds and evaluate on one, rotating across folds. We report mean  $\pm$  SD across folds for accuracy, macro-F1, weighted-F1, and per-class precision/recall/F1.

Hyperparameters are fixed across models; no tuning is performed on the external sets. At inference, we average fold-wise probabilities (simple ensemble) to stabilise predictions and optionally surface predictive entropy as an uncertainty signal for educator-facing use (not used to optimise metrics).

For summary metrics and per-class F1, we compute non-parametric bootstrap 95% confidence intervals (2,000–5,000 resamples, i.i.d.). When a video ID is available, we additionally provide cluster (block) bootstrap CIs by resampling videos to account for intra-video correlation.

#### 2.4. Statistical analysis

We treated the 10 cross-validation folds as repeated measurements for each model and computed macro-F1 scores per fold. This within-subjects design evaluates all models on the same data splits. Because fold-level scores ( $N = 10$ ) cannot be assumed to be normally distributed, we used non-parametric tests. A Friedman test assessed overall differences among models, followed by Wilcoxon signed-rank tests for pairwise comparisons.  $p$ -values were adjusted using the Holm method ( $\alpha = .05$ ). For each model, we report the mean macro-F1, standard deviation, and non-parametric bootstrap 95% confidence intervals. For pairwise comparisons, we additionally report mean differences ( $\Delta$ ) and paired-sample Cohen's  $d$ . Linear mixed-effects models were considered, but fold-level performance scores contain only ten observations per model, which limits the stability of variance-component estimation and leads to unreliable random-effects structures. Given this constraint, non-parametric repeated-measures tests offer a more appropriate and conservative alternative for evaluating model differences.

### 3. Results

#### 3.1. RQ1: Reliability of the Adapted IAM Coding Scheme

We draw on the original IAM (Gunawardena et al., 1997) and its adaptation for social media contexts by Nguyen and Diederich (2023) to develop a codebook suited to non-linear discourse environments (Table 1 and 2). Following Nguyen and Diederich (2023), we added the category non-knowledge construction (NonKC) to capture comments unrelated to learning processes. This addition is necessary because social media platforms contain large volumes of comments that are social, humorous, emotional, phatic, or otherwise not oriented toward epistemic engagement, and failing to account for these responses would bias the distribution of knowledge-construction categories and inflate model estimates. For knowledge-construction discourse, we retained all indicators from IAM Phases I and II and reorganised them into two categories: *Share* and *Explore*. Previous research has shown that higher-level IAM Phases IV and V rarely appear in online comments; removing or merging these phases into Phase III has been shown to produce more robust model performance (Gunawardena et al., 2023; Nguyen & Diederich, 2023). Accordingly, we combined these three phases into a single category, *Integrate*, and developed corresponding indicators based on the subcomponents specified in IAM Phases III–V. The final framework thus consists of four categories: NonKC, Share, Explore, and Integrate. This adaptation transforms IAM from a framework describing linear, iterative developmental phases into one that characterises discrete, non-sequential discourse categories at the comment level.

In the main annotated dataset of 20,000 comments, the application of this codebook achieved an inter-rater reliability of Cohen's  $\kappa = 0.79$ , indicating an almost perfect level of agreement (Landis & Koch, 1977). We then examined coder agreement in the complete random subsets from four external domains. Across these domains, Cohen's  $\kappa$  values ranged from .853 to .926, likewise falling within the “almost perfect” range. In addition to  $\kappa$ , annotation accuracy ranged from .913 to .958 and macro-F1 scores from .876 to .949 (see Table 8). These reliability estimates are based on a balanced sample of independently coded comments and do not adjust for clustering by video or domain; as such, they quantify agreement at the comment level rather than providing a cluster-adjusted  $\kappa$  for the natural, imbalanced distributions. This demonstrates that the codebook provides stable and reliable annotations across domains with markedly different discourse characteristics, indicating that IAM can indeed be adapted for knowledge construction assessment in noninteractive comments. These findings lay the foundation for the supervised training and evaluation of the automatic classifier reported in RQ2, as well as the cross-domain generalisation analyses presented in RQ3.

**Table 1**  
Knowledge construction categories (Part 1).

Gunawardena et al. (1997)	Nguyen & Diederich (2023)	Category Indicator
N/A	<b>Non-knowledge construction (NonKC)</b> Comment to socialise (positive and negative reactions), with less focus on the video's content. This code is added to capture sentiments.	<b>NonKC</b> <ul style="list-style-type: none"> <li>Express emotion</li> <li>Socialise or acknowledge others</li> <li>Produce irrelevant or incoherent content</li> <li>Ask unrelated questions to the learning content</li> </ul>
<b>PhI: Sharing/comparing of information</b> PhI/A. A statement of observation or opinion PhI/B. A statement of agreement from one or more other participants PhI/C. Corroborating examples provided by one or more participants PhI/D. Asking and answering questions to clarify details of statements PhI/E. Definition, description, or identification of a problem	<b>Share ideas</b> Ask clarifying questions, seek information or provide simple statements (personal experiences, facts or opinions). Comments are related to video content.	<b>Share</b> <ul style="list-style-type: none"> <li>State an opinion or observation (PhI/A)</li> <li>Express simple agreement (PhI/B)</li> <li>Provide simple example or information (PhI/C)</li> <li>Ask for clarification (PhI/D)</li> <li>Identify a problem or issue (PhI/E)</li> </ul>
<b>PhII: The discovery and exploration of dissonance or inconsistency among ideas, concepts or statements</b> PhII/A. Identifying and stating areas of disagreement PhII/B. Asking and answering questions to clarify the source and extent of disagreement PhII/C. Restating the participant's position, and possibly advancing arguments or considerations in its support by references to the participant's experience, literature, formal data collected, or proposal of relevant metaphor or analogy to illustrate point of view	<b>Explore dissonances</b> State agreement or disagreement (including simple statements such as 'I agree'). Ask questions to clarify the extent of disagreement.	<b>Explore</b> <ul style="list-style-type: none"> <li>Express disagreement (PhII/A)</li> <li>Question or probe another statement (PhII/B)</li> <li>Defend one's position with minimal support (PhII/C)</li> </ul>

### 3.2. RQ2: Comparison of Classifiers for IAM-Based Knowledge Construction

#### 3.2.1. Model performance and comparisons

Table 3 summarises the overall performance of all models in terms of mean accuracy, macro-F1 and weighted-F1 over 10-fold cross-validation. The two TF-IDF baselines show substantially lower performance (macro-F1  $\approx$  0.71–0.73) than all transformer-based models (macro-F1  $\approx$  0.81–0.84). This pattern indicates that, for our dataset of student contributions, pre-trained language models capture relevant linguistic and contextual cues more effectively than traditional bag-of-words approaches.

Among the transformer-based models, the DeBERTa variants achieve the best results and are tightly clustered in a narrow performance band. Our best configuration, *DeBERTa-v3-large* (ours: +Focal +LS +R-Drop), attains the highest mean macro-F1 (0.841  $\pm$  0.010), with a bootstrap 95% confidence interval of approximately [0.835, 0.848]. The ablation variants, where we removed R-Drop, label smoothing, or focal loss, show only small decreases in macro-F1 (on the order of 0.005–0.008), and *DeBERTa-v3-base* is

also close to the best configuration. Although we highlight the full DeBERTa-v3-large configuration as the best model in terms of mean accuracy and macro-F1, the performance differences among DeBERTa variants are modest.

Figure 3 visualises the distribution of 10-fold macro-F1 scores for each model. The TF-IDF baselines exhibit consistently lower scores and little overlap with the transformer models. In contrast, the transformer models form a compact group of high-performing models, with the DeBERTa variants showing both high central performance and relatively low variability across folds. Figure 4 further illustrates that, on almost every fold, transformer models outperform the TF-IDF baselines when evaluated on the same data splits, reinforcing the repeated-measures design of our analysis.

The Friedman test on fold-level macro-F1 scores across all models revealed a highly significant overall effect,  $\chi^2(9) = 77.38$ ,  $p < 10^{-12}$ , indicating systematic differences in performance. We therefore conducted pairwise Wilcoxon signed-rank tests with Holm correction. All transformer-based models significantly outperformed both TF-IDF baselines after correction for multiple comparisons. For example,

**Table 2**  
Knowledge construction categories (Part 2).

Gunawardena et al. (1997)	Nguyen & Diederich (2023)	Category Indicator
<p><b>PhIII: Negotiation of meaning/co-construction of knowledge</b> PhIII/A. Negotiation or clarification of the meaning of terms PhIII/B. Negotiation of the relative weight to be assigned to types of argument PhIII/C. Identification of areas of agreement or overlap among conflicting concepts PhIII/D. Proposal and negotiation of new statements embodying compromise, co-construction PhIII/E. Proposal of integrating or accommodating metaphors or analogies</p>	<p><b>Negotiate</b> Clarify concepts. Propose and negotiate areas of disagreement to integrate ideas. Employ more extensive evidence and explanation than earlier phases.</p>	<p><b>Integrate</b></p> <ul style="list-style-type: none"> <li>• Clarify concepts or terms (PhIII/A)</li> <li>• Propose or synthesize integrative statements (PhIII/C, PhIII/D, PhIII/E, PhV/A)</li> <li>• Evaluate statements by comparing perspectives, experiences, or information sources (PhIII/B, PhIV/A, PhIV/B, PhIV/C, PhIV/D, PhIV/E)</li> <li>• Apply or reflect on new understanding (PhV/B, PhV/C)</li> </ul>
<p><b>PhIV: Testing and modification of proposed synthesis or co-construction</b> PhIV/A. Testing the proposed synthesis against "received fact" as shared by the participants and/or their culture PhIV/B. Testing against existing cognitive schema PhIV/C. Testing against personal experience PhIV/D. Testing against formal data collected PhIV/E. Testing against contradictory testimony in the literature</p>	<p><b>Test/modify proposed ideas</b> Test the proposed idea syntheses against other contexts (data, references, personal experiences).</p>	
<p><b>PhV: Agreement statement(s)/applications of newly-constructed meaning</b> PhV/A. Summarization of agreement(s) PhV/B. Applications of new knowledge PhV/C. Metacognitive statements by the participants illustrating their understanding that their knowledge or ways of thinking (cognitive schema) have changed as a result of the conference interaction</p>	<p><b>Apply new knowledge</b> Summarise agreement, make reflective statements and apply new knowledge.</p>	

**Table 3**  
Overall performance (10-fold CV). Values are mean  $\pm$  SD. The best model (based on mean Accuracy) is shown in bold.

Model	Accuracy	Macro-F1	Weighted-F1
TF-IDF + Logistic Regression	.729 $\pm$ .012	.728 $\pm$ .012	.728 $\pm$ .012
TF-IDF + LinearSVC	.711 $\pm$ .009	.709 $\pm$ .009	.709 $\pm$ .009
DistilBERT (base)	.807 $\pm$ .010	.807 $\pm$ .010	.807 $\pm$ .010
RoBERTa-base	.821 $\pm$ .010	.821 $\pm$ .009	.821 $\pm$ .009
DeBERTa-v3-base	.834 $\pm$ .007	.834 $\pm$ .007	.834 $\pm$ .007
DeBERTa-v3-large (base backbone)	.833 $\pm$ .010	.834 $\pm$ .010	.834 $\pm$ .010
<b>DeBERTa-v3-large (ours: +Focal +LS +R-Drop)</b>	<b>.841 <math>\pm</math> .011</b>	<b>.841 <math>\pm</math> .010</b>	<b>.841 <math>\pm</math> .010</b>
- no R-Drop	.833 $\pm$ .011	.834 $\pm$ .010	.834 $\pm$ .010
- no LS	.840 $\pm$ .009	.841 $\pm$ .009	.841 $\pm$ .009
- no Focal	.836 $\pm$ .008	.836 $\pm$ .008	.836 $\pm$ .008

the contrast between *DeBERTa-v3-large (ours)* and TF-IDF + Logistic Regression yielded an average macro-F1 difference of  $\Delta \approx 0.113$  in favour of our model, with a Holm-adjusted  $p < .05$  and a large paired effect size. From an educational perspective, this means that transformer models are much more reliable in assigning the correct knowledge construction category to student contributions than the simpler TF-IDF approaches.

Differences among the transformer models themselves were smaller. Although *DeBERTa-v3-large (ours)* achieved slightly higher mean macro-F1 than *DeBERTa-v3-base* and *RoBERTa-base* (e.g.,  $\Delta \approx 0.008$  and  $\Delta \approx 0.021$ , respectively), these improvements did not remain statistically significant after Holm correction. Likewise, the ablated *DeBERTa-v3-large* variants (without R-Drop, label smoothing, or focal loss) did not differ significantly from the full model. In practical terms, our full *DeBERTa* configuration is numerically the best-performing model, but its advantage over other *DeBERTa* variants is modest; the most robust result is the large and statistically significant performance gap between any transformer-based model and the TF-IDF baselines. Consequently, we adopt the full *DeBERTa-v3-large* configuration as our primary model for subsequent analyses, not because its performance differences are statistically decisive, but because it consistently attains the highest macro-F1 and exhibits the most stable and well-calibrated probability estimates across folds.

**Table 4**Performance comparison across models for the **NonKC** category.

Model	Precision	Recall	F1
TF-IDF + Logistic Regression	.734 ± .020	.736 ± .026	.735 ± .019
TF-IDF + LinearSVC	.714 ± .021	.725 ± .023	.719 ± .018
DistilBERT (base)	.831 ± .015	.782 ± .029	.805 ± .013
RoBERTa-base	.858 ± .024	.782 ± .044	.817 ± .018
DeBERTa-v3-base	<b>.883</b> ± .024	.782 ± .029	.829 ± .011
DeBERTa-v3-large (replaced with base backbone)	.878 ± .022	.790 ± .025	.831 ± .012
DeBERTa-v3-large (ours: +Focal +LS +R-Drop)	.861 ± .032	<b>.827</b> ± .038	.843 ± .017
- no R-Drop	.873 ± .022	.802 ± .031	.836 ± .013
- no LS	.869 ± .033	.826 ± .035	<b>.846</b> ± .009
- no Focal	.873 ± .033	.808 ± .036	.838 ± .012

### 3.2.2. Per-category performance

To better understand how the classifiers behave with respect to the different categories of knowledge construction, we analysed precision, recall, and F1 for each category separately (Tables 4–7). Across all four categories, transformer-based models consistently outperform the TF-IDF baselines, but the relative difficulty of the categories and the benefits of the DeBERTa variants differ slightly by category.

For the NonKC category (Table 4), all models perform relatively well, which is expected given that nonKC comments often contain clearer superficial or off-topic cues. The TF-IDF baselines still lag behind (F1  $\approx$  .72–.74), but the gap to transformers is somewhat smaller than for the knowledge-building categories. Within the transformer family, DeBERTa-v3-base has the highest precision (.883 ± .024), whereas the “no LS” variant attains the highest F1 (.846 ± .009). Our full model strikes a balance, with slightly lower precision (.861 ± .032) but the highest recall across all models (.827 ± .038), yielding an F1 of .843 ± .017. This pattern suggests that the full configuration conservatively avoids misclassifying knowledge-construction comments as NonKC, which is advantageous when the model is used to identify meaningful knowledge-construction episodes.

The Share category (Table 5) is more challenging for all models and shows the lowest scores among the knowledge construction categories. The TF-IDF baselines perform markedly worse here (F1  $\approx$  .62–.65), whereas transformer models yield F1 scores in the mid-.70s to high-.79s. Our full DeBERTa configuration achieves an F1 of .799 ± .014 with balanced precision (.796 ± .029) and recall (.804 ± .030). The “no LS” variant obtains a very similar F1 (.797 ± .016) with slightly higher precision but slightly lower recall. DeBERTa-v3-base yields the highest recall for Share (.817 ± .024), suggesting that label smoothing may help the model recover more borderline sharing contributions at the cost of a small precision loss. These results indicate that distinguishing Share from the other categories is inherently challenging at both the conceptual and linguistic levels, which may partly reflect the fact that the operational definition of Share in this study encompasses a broader and more diverse set of behaviors than the other categories.

For the Explore category (Table 6), performance again improves substantially when moving from TF-IDF (F1  $\approx$  .66–.69) to transformer-based models (F1  $\approx$  .77–.82). Among transformers, the DeBERTa variants again dominate. Our full model achieves the highest F1 (.821 ± .016), with the

**Table 5**Performance comparison across models for the **Share** category.

Model	Precision	Recall	F1
TF-IDF + Logistic Regression	.652 ± .021	.646 ± .023	.648 ± .017
TF-IDF + LinearSVC	.634 ± .016	.614 ± .027	.623 ± .018
DistilBERT (base)	.752 ± .022	.782 ± .032	.766 ± .017
RoBERTa-base	.768 ± .025	.785 ± .027	.775 ± .011
DeBERTa-v3-base	.772 ± .015	<b>.817</b> ± .024	.794 ± .009
DeBERTa-v3-large (replaced with base backbone)	.781 ± .022	.800 ± .021	.790 ± .013
DeBERTa-v3-large (ours: +Focal +LS +R-Drop)	<b>.796</b> ± .029	.804 ± .030	<b>.799</b> ± .014
- no R-Drop	.777 ± .027	.804 ± .037	.789 ± .013
- no LS	.793 ± .032	.803 ± .049	.797 ± .016
- no Focal	.775 ± .020	.812 ± .020	.793 ± .009

**Table 6**Performance comparison across models for the **Explore** category.

Model	Precision	Recall	F1
TF-IDF + Logistic Regression	.693 ± .022	.679 ± .029	.686 ± .023
TF-IDF + LinearSVC	.679 ± .021	.650 ± .022	.664 ± .018
DistilBERT (base)	.771 ± .034	.776 ± .041	.772 ± .016
RoBERTa-base	.778 ± .021	.816 ± .028	.796 ± .016
DeBERTa-v3-base	.789 ± .024	.836 ± .029	.812 ± .016
DeBERTa-v3-large (replaced with base backbone)	.781 ± .033	.845 ± .028	.811 ± .016
DeBERTa-v3-large (ours: +Focal +LS +R-Drop)	<b>.803</b> ± .032	.842 ± .026	<b>.821</b> ± .016
- no R-Drop	.785 ± .034	.841 ± .023	.812 ± .014
- no LS	.794 ± .025	<b>.850</b> ± .030	.820 ± .015
- no Focal	.802 ± .026	.834 ± .036	.817 ± .015

**Table 7**Performance comparison across models for the **Integrate** category.

Model	Precision	Recall	F1
TF-IDF + Logistic Regression	.833 ± .015	.855 ± .013	.844 ± .011
TF-IDF + LinearSVC	.808 ± .011	.854 ± .015	.830 ± .009
DistilBERT (base)	.885 ± .028	.890 ± .028	.887 ± .013
RoBERTa-base	.890 ± .027	<b>.902</b> ± .036	.895 ± .014
DeBERTa-v3-base	.906 ± .031	.899 ± .030	.902 ± .013
DeBERTa-v3-large (replaced with base backbone)	.909 ± .025	.899 ± .034	<b>.903</b> ± .013
DeBERTa-v3-large (ours: +Focal +LS +R-Drop)	.917 ± .014	.891 ± .031	<b>.903</b> ± .012
- no R-Drop	.916 ± .024	.885 ± .035	.899 ± .013
- no LS	<b>.920</b> ± .023	.883 ± .026	.900 ± .006
- no Focal	.907 ± .033	.888 ± .040	.897 ± .012

best precision (.803 ± .032) and strong recall (.842 ± .026). The “no LS” variant achieves the highest recall (.850 ± .030) and an F1 of .820 ± .015, essentially matching the full model. These patterns suggest that the additional regularisation from focal loss, label smoothing, and R-Drop helps DeBERTa-v3-large to capture exploratory moves slightly more reliably, without sacrificing precision.

For the Integrate category (Table 7), all transformer models achieve very high performance, with F1 scores around or above .89. The TF-IDF baselines lag behind (F1  $\approx$  .83–.84), indicating that capturing negotiation moves requires richer contextual representations than simple bag-of-words features can provide. Within the transformer family, the DeBERTa variants obtain the strongest results. Our full configuration (*DeBERTa-v3-large*, +Focal +LS +R-Drop) attains an F1 of .903 ± .012, matching the best-performing large-backbone variant and slightly exceeding DeBERTa-v3-base (F1 = .902 ± .013). The “no LS” ablation achieves the highest precision (.920 ± .023) for Integrate but trades this for somewhat lower recall (.883 ± .026), suggesting a more conservative classification that misses some genuine contributions.

In summary, transformer-based models deliver consistent improvements across all knowledge-construction categories and clearly outperform TF-IDF baselines, with the DeBERTa-v3-large variants performing best. Share and

**Table 8**

Overall performance of *DeBERTa-v3-large (ours)* and human agreement on the random subsets. Values are computed against coder B (for the model) and between coder B and coder C (for human agreement).

Domain	<i>N</i>	Acc (model vs B)	Macro-F1 (model vs B)	Weighted-F1 (model vs B)	Acc (B vs C)	Macro-F1 (B vs C)	$\kappa$ (B,C)
Medicine	600	.897	.864	.896	.958	.949	.926
Language	600	.820	.707	.803	.940	.893	.865
Music	600	.752	.705	.738	.948	.923	.913
Programming	600	.882	.818	.881	.913	.876	.853

Explore remain the most difficult to classify due to their conceptual overlap and linguistic variability, whereas NonKC and Integrate are comparatively clearer. Although performance differences among the DeBERTa configurations are small and not statistically significant, the full model shows the highest and most stable results. These findings indicate that modern transformer architectures can provide sufficiently accurate comment-level category labels for large-scale analyses of students' knowledge-construction behaviors.

### 3.3. RQ3: Cross-Domain Generalisation of DeBERTa-KC

#### 3.3.1. Overall cross-domain performance on realistic samples

Figure 5 shows distribution of knowledge-construction categories in the Normal subset across four domains. In all domains, *NonKC* is the most frequent category, followed by *Share*, whereas *Explore* and *Integrate* occur far less often. In the language domain, *NonKC* is dominant (71%), *Share* accounts for a moderate proportion (24.8%), and both *Explore* (2.7%) and *Integrate* (1.8%) appear only rarely. In the medicine domain, *NonKC* (47.8%) remains slightly more prevalent than *Share* (45.2%), while *Explore* (3.0%) and *Integrate* (4.0%) show limited but noticeable presence. In the music domain, *NonKC* has the largest share (48.2%), followed by *Share* (40.7%); *Explore* (3.8%) and *Integrate* (7.3%) occur somewhat more frequently than in the medicine domain. Finally, in the programming domain, *NonKC* continues to be the largest category (57.5%), *Share* remains substantial (27.7%), and *Explore* (5.7%) and *Integrate* (9.2%) reach their highest proportions across all domains. This pronounced skew, particularly in the language domain, also helps explain the lower macro-F1 scores observed there: domains with very few *Explore* and *Integrate* instances provide limited training and evaluation data for these stages, which depresses their per-category F1 and the overall macro-F1.

Table 8 summarises overall classifier performance in each domain on the Normal subsets (600 comments per domain), using coder B as the reference. Human agreement (coder B vs. coder C) on these same subsets is also shown as an approximate upper bound.

Compared to the in-domain macro-F1 of approximately .841 reported in Section 3.2, the classifier's macro-F1 on the Normal subsets is slightly higher in *Medicine* (.864) and somewhat lower in *Programming* domain (.818), yielding small absolute differences of +0.023 and -0.023, respectively. In contrast, performance drops more substantially in *Language* and *Music* domains, with macro-F1 scores of

.707 and .705, corresponding to reductions of approximately -0.13 relative to the in-domain result. Part of this drop can be attributed to the natural distribution of categories in those domains, such as *Language*, which contain very few examples of *Explore* and *Integrate*, disproportionately affecting the macro-F1. In domains where higher-order epistemic moves are almost nonexistent, even small absolute errors translate into large proportional reductions in macro-level metrics. The weighted-F1 scores follow a similar pattern, reflecting the strong influence of the abundant *NonKC* and *Share* categories.

Figure 6 visualises these macro-F1 scores alongside the human-human macro-F1, with a dashed line indicating the in-domain macro-F1. The figure highlights two key points. First, in all domains there remains a sizeable gap between model-human agreement and human-human agreement (typically around .08-.14 macro-F1), indicating that the classifier has not yet reached human reliability. Second, cross-domain degradation is not uniform: the model retains most of its in-domain performance in *Medicine* and *Programming* domains, but exhibits noticeably weaker performance in *Language* and *Music* domains.

To examine whether classifier performance varied across domains, we fitted a logistic regression model in which the binary outcome indicated whether the model's prediction matched Coder B's annotation (correct vs. incorrect), with domain as a categorical predictor. A likelihood ratio test showed a significant domain effect ( $\chi^2(3) = 56.81, p < .001$ ), indicating that the probability of correct prediction depended on the educational domain. Using *Language* domain as the reference category, the classifier was more likely to be correct in *Medicine* ( $\approx 1.90$ ) and *Programming* ( $\approx 1.64$ ) domains, but less likely in *Music* ( $\approx 0.66$ ) domain.

#### 3.3.2. Per-category cross-domain generalisation

To further examine which categories generalize well across domains and which remain challenging, we conducted a category-level analysis using the Normal and Augment subsets. Because the Augment subset oversampled previously rare stages, the resulting evaluation set does not reflect the natural class distribution but provides a more balanced basis for assessing the classifier's ability to identify each category across domains. Using all available comments (Normal + Augment), we computed precision, recall, and F1 for each category relative to Coder B's annotations (see Table 9 and Figure 8).

The results reveal several consistent patterns. *NonKC* was the easiest category to identify in all domains, with F1 ranging from .828 to .954. *Integrate* also showed robust

**Table 9**

Per-category F1 scores by domain (all samples: Random + Augment), computed against coder B's labels.

Category	Medicine	Language	Music	Programming
Explore	.746	.733	.693	.748
Integrate	.826	.779	.809	.824
Share	.778	.455	.543	.661
NonKC	.923	.891	.828	.954

performance across domains, with F1 between .779 and .826, consistent with the within-domain results reported in Section 3.2.2, where *NonKC* and *Integrate* were likewise the best-performing categories. *Explore* showed moderate yet stable performance, with F1 ranging from .693 to .748, indicating that the classifier captured exploratory behaviours with reasonable consistency. In contrast, *Share* exhibited the greatest cross-domain variability and was clearly the most fragile category. Performance was relatively stronger in the Medicine (F1 = .778) and Programming (F1 = .661) domains, but dropped substantially in Language (F1 = .455) and Music (F1 = .543).

To better understand the error patterns behind these scores, we examined the row-normalised confusion matrices in Figure 7. Across all four domains, most *Explore* and *Integrate* comments lie on the diagonal (typically 90–97% correct for *Explore* and 92–100% for *Integrate*), with the main off-diagonal mass corresponding to mutual confusion between these two adjacent stages rather than with *NonKC*. *NonKC* is also clearly separated: 94–98% of gold *NonKC* comments are correctly classified, and misclassifications into knowledge-construction categories are rare and mostly routed through *Share*. In contrast, errors are concentrated in the *Share* rows. In Medicine and Programming, many *Share* comments are redistributed to *Explore* or *Integrate*, whereas in Language and Music a substantial proportion of *Share* comments are instead predicted as *NonKC* (e.g., 35% in Language and 30% in Music). This pattern likely reflects that certain domain-specific behaviours, such as sentence construction in language learning or expressive performance-related actions in music learning, were coded by humans as *Share* but classified by the model as *NonKC* because such behaviours were absent from the training data. It may also be influenced by the fact that *Share* is the category with the lowest overall classification performance, partly because it encompasses a broader and more heterogeneous set of behaviours than the other categories.

In summary, the results indicate that the classifier can generalise across educational domains, but with limitations. At the aggregate level, the model retained much of its within-domain performance in the Medicine and Programming domains, whereas accuracy decreased substantially in the Language and Music domains. At the category level, the transformer model's strengths were not confined to a single category, *NonKC*, *Explore*, and *Integrate* all achieved moderate to substantial F1 scores across domains. The primary cross-domain weakness was *Share*, particularly in domains

with strong domain-specific discourse features, suggesting a need for domain-sensitive calibration or additional training data.

## 4. Discussion

This study developed and validated a framework combining a theoretically grounded codebook with a high-performing transformer classifier to analyse knowledge construction (KC) in non-linear, loosely connected online discourse. Our findings demonstrate that adapting interaction-oriented models like the IAM to comment-level classification is feasible and informative at scale.

The revised four-category IAM scheme preserves core conceptual distinctions (Gunawardena et al., 1997) while addressing limitations in unstructured environments (Dubovi & Tabak, 2020; Wever et al., 2006). By consolidating higher-level phases into *Integrate* and incorporating a *NonKC* category (Nguyen & Diederich, 2023), the scheme trades some granularity for coding reliability, achieving performance comparable to or exceeding earlier adaptations (Gunawardena et al., 2023; Nguyen & Diederich, 2023).

Unlike prior approaches dependent on sequential interactions, our framework detects epistemic moves without requiring thread continuity. Cross-domain evaluation shows that *Integrate* generalises well, while *Share* exhibits greater linguistic variability, indicating that category design influences transferability. This positions the adapted IAM as suitable for large-scale, comment-based environments such as YouTube.

Transformer models, particularly DeBERTa variants, outperformed traditional TF-IDF approaches, aligning with recent NLP research (Ba et al., 2023; Castellanos-Reyes et al., 2025; He, Liu, et al., 2021). Ablation analyses suggest that additional techniques like label smoothing mainly improve stability and calibration rather than accuracy, indicating that well-configured base transformers suffice for many KC tasks. The model's ability to generalise *Explore* and *Integrate* across domains suggests it captures deeper epistemic functions beyond surface lexical patterns, whereas *Share* remains sensitive to pragmatic variation. These results also connect back to the construct-validity concerns discussed in Section 1.3. The fact that *Explore* and *Integrate* generalised more consistently across domains than might be expected from surface lexical patterns indicates that the model is not merely learning topic-specific cues, but is partially capturing deeper epistemic functions. Conversely, the instability of *Share* aligns with the prediction that heterogeneous, pragmatically variable categories are more vulnerable to lexical approximation. This provides empirical evidence that automated IAM coding can approximate epistemic constructs, but only when those constructs have sufficiently distinct linguistic realisations.

Cross-domain evaluation provides further insight into the strengths and limits of automated KC classification. Similar to findings in domain adaptation research (Gururangan et al., 2020; Peters et al., 2019), our results show that models trained in one domain can generalise to others, but performance varies depending on the linguistic and

pragmatic characteristics of the target domain. In structured and content-heavy domains such as medicine and programming, Explore and Integrate categories remained relatively stable, echoing observations that Transformers are adept at capturing domain-general semantic patterns (He, Liu, et al., 2021). However, the reduced performance on Share in language- and music-related channels mirrors prior reports that low-elaboration contributions often rely on subtle pragmatic markers (Nguyen & Diederich, 2023) that differ across disciplinary and community contexts. These patterns indicate that some aspects of knowledge construction are more domain-sensitive than others, and they underline the value of light domain adaptation or selective human review for categories with high linguistic variability.

Overall, this study shows that IAM-based coding can be reliably adapted for large, non-interactive comment streams and that modern transformers can support semi-automated KC analysis at unprecedented scale and performance in informal online learning environments.

#### 4.1. Implications for research and practice

This study demonstrates a practical approach for adapting the IAM framework to large-scale, nonlinear, noninteractive datasets. By redefining the phases as comment-level categories and incorporating an explicit NonKC label, we preserved meaningful theoretical distinctions while producing a schema that is more feasible for annotation and modeling. This approach enables researchers to conduct reliable content analyses in large, fragmented online learning environments and supports practitioners in efficiently distinguishing learning-related from non-learning contributions.

The results also indicate that transformer-based models can effectively support, rather than replace, human discourse analysis. Although DeBERTa-KC does not reach the level of human-human agreement, its performance is sufficient for generating aggregate indicators and enabling semi-automated workflows, and it clearly surpasses traditional machine-learning baselines. Prediction entropy further provides a useful mechanism for identifying low-confidence outputs that require human review.

Cross-domain analyses suggest that our classifier generalises beyond its training context, particularly in domains with more structured and task-oriented discourse. The model provides reliable signals for NonKC, Explore, and Integrate, whereas Share remains the most sensitive to domain-specific linguistic variation. This behaviour not only highlights the need for light domain calibration or selective human oversight when deploying automated KC classification, but also offers insight into the IAM categories themselves. Categories with clearer conceptual boundaries, such as Explore and Integrate, were consistently recoverable across domains, while Share, operationalised as a broad and heterogeneous category, showed substantially weaker stability. These patterns suggest that future adaptations of the IAM for large and unstructured environments may benefit from refining or subdividing Share to reduce its pragmatic variability.

#### 4.2. Limitations and directions for future work

First, our reliability analyses treat comments as independent units and do not compute cluster-adjusted agreement coefficients by video or domain, which may differ somewhat from reliability estimates obtained on fully natural, imbalanced samples. Although we adapted IAM to four broad categories and demonstrated high reliability, the scheme still represents a simplification of the original five-phase model and of the rich spectrum of online epistemic moves. In particular, the combined *Integrate* category conflates activities such as testing, integration, and application. Future work could explore whether more fine-grained subcategories within *Integrate* can be recovered reliably in specific domains or with additional contextual information.

Second, our modeling approach treats each comment as an isolated unit. However, some platforms support threaded interactions, enabling comments to unfold within conversational structures. Incorporating sequence models that account for thread structure or temporal dynamics may better capture how knowledge construction develops over time.

Third, although we examined cross-domain generalisation on four additional YouTube channels, all data in this study come from a single platform and language context. Applying the approach to different platforms and other languages or cultural settings would provide a stronger test of generalisability and might require further adaptation of the coding scheme.

Fourth, the external validation domains, while diverse, are not fully representative of all educational uses of social media. For instance, we did not include domains with highly contentious or polarised debates, where epistemic and socio-political dimensions are deeply intertwined, nor did we examine settings with strong teacher presence and scaffolding. Investigating how IAM categories manifest in such contexts could refine both the coding scheme and model behaviour.

Finally, we did not explore more advanced forms of model adaptation or fairness-aware training. Domain adaptation techniques (e.g., adversarial alignment, parameter-efficient fine-tuning) might improve performance in challenging domains without requiring extensive re-annotation. Likewise, future work should examine whether model errors are systematically patterned by learner characteristics (where available) or by topic, and consider bias-sensitive evaluation and mitigation strategies.

### 5. Conclusion

This study proposed and validated a framework for analyzing knowledge construction within nonlinear and unstructured text in online learning environments. By adapting four categories of knowledge construction from the Interaction Analysis Model, we developed a codebook suited to these environments, including NonKC, Share, Explore, and Integrate. Building on this codebook, we introduced DeBERTa-KC, an automated classifier based on the DeBERTa-V3-large architecture, which outperformed baseline models and showed stable generalization in external validations. The

categories of NonKC, Explore, and Integrate transferred more consistently across contexts, whereas Share appeared more context dependent. Overall, this study provides an empirically validated and scalable approach for identifying knowledge construction in unstructured learner discourse. The findings open new opportunities for systematic and large-scale investigations of learners' cognitive processes in online learning environments and offer meaningful implications for learning analytics and future research.

## References

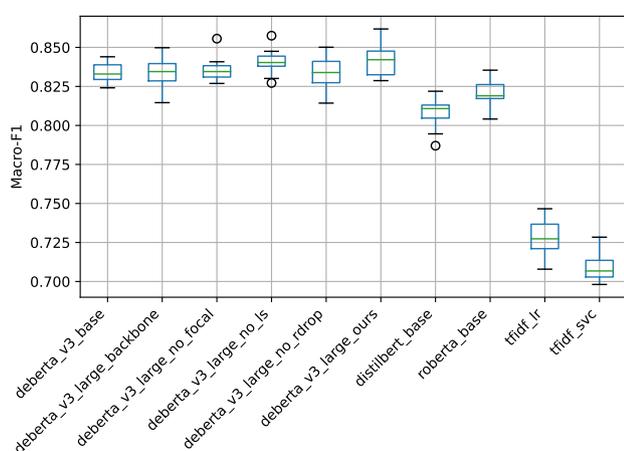
- Aragón, P., Gómez, V., García, D., & Kaltenbrunner, A. (2017). Generative models of online discussion threads: State of the art and research challenges. *Journal of Internet Services and Applications*, 8, 15. <https://doi.org/10.1186/s13174-017-0066-z>
- Ba, S., Hu, X., Stein, D., & Liu, Q. (2023). Assessing cognitive presence in online inquiry-based discussion through text classification and epistemic network analysis. *British Journal of Educational Technology*, 54. <https://doi.org/10.1111/bjet.13285>
- Breivik, J. (2016). Critical thinking in online educational discussions measured as progress through inquiry phases: A discussion of the cognitive presence construct in the community of inquiry framework. *International Journal of E-Learning & Distance Education*, 32(1), 1–16.
- Castellanos-Reyes, D., Olesova, L., & Sadaf, A. (2025). Transforming online learning research: Leveraging gpt large language models for automated content analysis of cognitive presence. *Internet and Higher Education*, 65. <https://doi.org/10.1016/j.iheduc.2025.101001>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT*, 4171–4186. <https://aclanthology.org/N19-1423/>
- Dornauer, V., Netzer, M., Kaczkó, É., Norz, L. M., & Ammenwerth, E. (2024). Automatic classification of online discussions and other learning traces to detect cognitive presence. *International Journal of Artificial Intelligence in Education*, 34. <https://doi.org/10.1007/s40593-023-00335-4>
- Dubovi, I., & Tabak, I. (2020). An empirical analysis of knowledge co-construction in youtube comments. *Computers and Education*, 156. <https://doi.org/10.1016/j.compedu.2020.103939>
- Ferreira, M., Neto, V., Rolim, V., Barbosa, A., Lins, R. D., Chen, G., & Gašević, D. (2020). Towards automatic content analysis of social presence in transcripts of online discussions. *Proceedings of the 10th International Learning Analytics & Knowledge Conference (LAK 2020)*, 329–338. <https://doi.org/10.1145/3375462.3375495>
- Garrison, D. R., Anderson, T., & Archer, W. (2001). Critical thinking, cognitive presence, and computer conferencing in distance education. *The American Journal of Distance Education*, 15(1), 7–23. <https://doi.org/10.1080/08923640109527071>
- Gašević, D., Dawson, S., & Siemens, G. (2015). Let's not forget: Learning analytics are about learning. *TechTrends*, 59(1), 64–71. <https://doi.org/10.1007/s11528-014-0822-x>
- Greenhow, C., & Lewin, C. (2016). Social media and education: Reconceptualizing the boundaries of formal and informal learning. *Learning, Media and Technology*, 41. <https://doi.org/10.1080/17439884.2015.1064954>
- Gunawardena, C. N., Chen, Y., Flor, N., & Sánchez, D. (2023). Deep learning models for analyzing social construction of knowledge online. *Online Learning*, 27. <https://doi.org/10.24059/olj.v27i4.4055>
- Gunawardena, C. N., Lowe, C. A., & Anderson, T. (1997). Analysis of a global online debate and the development of an interaction analysis model for examining social construction of knowledge in computer conferencing. *Journal of Educational Computing Research*, 17(4), 395–429. <https://doi.org/10.2190/7MQV-X9UJ-C7Q3-NRAG>
- Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020). Don't stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964*. <https://doi.org/10.48550/arXiv.2004.10964>
- Hallgren, K. A. (2012). Computing inter-rater reliability for observational data: An overview and tutorial. *Tutorials in Quantitative Methods for Psychology*, 8. <https://doi.org/10.20982/tqmp.08.1.p023>
- He, P., Gao, J., & Chen, W. (2021). Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*. <https://doi.org/10.48550/arXiv.2111.09543>
- He, P., Liu, X., Gao, J., & Chen, W. (2021). Deberta: Decoding-enhanced bert with disentangled attention. *ICLR 2021 - 9th International Conference on Learning Representations*. <https://openreview.net/forum?id=XPZLaotutsD>
- Henri, F. (1992). Computer conferencing and content analysis. In A. R. Kaye (Ed.), *Collaborative learning through computer conferencing* (pp. 117–136). Springer Berlin Heidelberg. [https://doi.org/10.1007/978-3-642-77684-7\\_8](https://doi.org/10.1007/978-3-642-77684-7_8)
- Herring, S. C. (1999). Interactional coherence in cmc. *Journal of Computer-Mediated Communication*, 4. <https://doi.org/10.1111/j.1083-6101.1999.tb00106.x>
- Hew, K. F., & Cheung, W. S. (2011). Higher-level knowledge construction in asynchronous online discussions: An analysis of group size, duration of online discussion, and student facilitation techniques. *Instructional Science*, 39(3), 303–319. <https://doi.org/10.1007/s11251-010-9129-2>
- Hmelo-Silver, C. E. (2003). Analyzing collaborative knowledge construction: Multiple methods for integrated understanding. *Computers and Education*, 41. <https://doi.org/10.1016/j.compedu.2003.07.001>
- Hu, Y., Donald, C., & Giacaman, N. (2022). A revised application of cognitive presence automatic classifiers for moocs: A new set of indicators revealed? *International Journal of Educational Technology in Higher Education*, 19. <https://doi.org/10.1186/s41239-022-00353-7>
- Hu, Y., Donald, C., & Giacaman, N. (2023). Can multi-label classifiers help identify subjectivity? a deep learning approach to classifying cognitive presence in moocs. *International Journal of Artificial Intelligence in Education*, 33. <https://doi.org/10.1007/s40593-022-00310-5>
- Hu, Y., Mello, R. F., & Gašević, D. (2021). Automatic analysis of cognitive presence in online discussions: An approach using deep learning and explainable artificial intelligence. *Computers and Education: Artificial Intelligence*, 2. <https://doi.org/10.1016/j.caeai.2021.100037>
- Kimmerle, J., Moskaliuk, J., Oeberst, A., & Cress, U. (2015). Learning and collective knowledge construction with social media: A process-oriented perspective. *Educational Psychologist*, 50. <https://doi.org/10.1080/00461520.2015.1036273>
- Kovanović, V., Joksimović, S., Waters, Z., Gašević, D., Kitto, K., Hatala, M., & Siemens, G. (2016). Towards automated content analysis of discussion transcripts: A cognitive presence case. *Proceedings of the 6th International Conference on Learning Analytics & Knowledge (LAK 2016)*, 15–24. <https://doi.org/10.1145/2883851.2883950>
- Krykoniuk, K., Hopkin-King, C., & Roberts, S. G. (2025). Developing a discourse space for analysing online discourse. *Discourse, Context & Media*, 67, 100929. <https://doi.org/10.1016/j.dcm.2025.100929>
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174. <https://doi.org/10.2307/2529310>
- Liang, X., Wu, L., Li, J., Wang, Y., Meng, Q., Qin, T., Chen, W., Zhang, M., & Liu, T. Y. (2021). R-drop: Regularized dropout for neural networks [Also available at arXiv:2106.14448]. *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*. <https://papers.nips.cc/paper/2021/hash/5a66b9200f29ac3fa0ae244cc2a51b39-Abstract.html>

- Lin, T. Y., Goyal, P., Girshick, R., He, K., & Dollar, P. (2020). Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42. <https://doi.org/10.1109/TPAMI.2018.2858826>
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*. <https://doi.org/10.48550/arXiv.1907.11692>
- Lucas, M., Gunawardena, C., & Moreira, A. (2014). Assessing social construction of knowledge online: A critique of the interaction analysis model. *Computers in Human Behavior*, 30, 574–582. <https://doi.org/10.1016/j.chb.2013.07.050>
- Marcoccia, M. (2004). On-line polylogues: Conversation structure and participation framework in internet newsgroups. *Journal of Pragmatics*, 36, 115–145. [https://doi.org/10.1016/S0378-2166\(03\)00038-9](https://doi.org/10.1016/S0378-2166(03)00038-9)
- Mayer, R. E. (1996). Learning strategies for making sense out of expository text: The soi model for guiding three cognitive processes in knowledge construction. *Educational Psychology Review*, 8(4), 357–371. <https://doi.org/10.1007/BF01463939>
- Mim, F. S., Aeron, S., Miller, E., & Wendell, K. (2025). From words to wisdom: Discourse annotation and baseline models for student dialogue understanding. *arXiv preprint arXiv:2511.20547*. <https://doi.org/10.48550/arXiv.2511.20547>
- Naab, T. K., Rueß, H., & Küchler, C. (2023). The influence of the deliberative quality of user comments on the number and quality of their reply comments. *New Media & Society*, 27, 62–83. <https://doi.org/10.1177/14614448231172168>
- Neto, V., Rolim, V., Pinheiro, A., Lins, R. D., Gašević, D., & Ferreira Mello, R. (2021). Automatic content analysis of online discussions for cognitive presence: A study of the generalizability across educational contexts. *IEEE Transactions on Learning Technologies*, 14(3), 299–312. <https://doi.org/10.1109/TLT.2021.3083178>
- Nguyen, H., & Diederich, M. (2023). Facilitating knowledge construction in informal learning: A study of tiktok scientific, educational videos. *Computers and Education*, 205. <https://doi.org/10.1016/j.compedu.2023.104896>
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in python. *The Journal of Machine Learning Research*, 12, 2825–2830. <https://jmlr.org/papers/v12/pedregosa11a.html>
- Peters, M. E., Ruder, S., & Smith, N. A. (2019). To tune or not to tune? adapting pretrained representations to diverse tasks. *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, 7–14. <https://doi.org/10.18653/v1/W19-4302>
- Sakeef, N., Dewan, M. A. A., Lin, F., & Parmar, D. (2025). Detecting cognitive engagement in online course forums: A review of frameworks and methodologies. *Natural Language Processing Journal*, 11(100146). <https://doi.org/10.1016/j.nlp.2025.100146>
- Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). Distilbert, a distilled version of bert: Smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*. <https://doi.org/10.48550/arXiv.1910.01108>
- Scardamalia, M., & Bereiter, C. (2006). Knowledge building: Theory, pedagogy, and technology. In R. K. Sawyer (Ed.), *Cambridge handbook of the learning sciences* (pp. 397–417). Cambridge University Press. <https://doi.org/10.1017/CBO9781139519526.025>
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/CVPR.2016.308>
- Tanskanen, S.-K. (2021). Fragmented but coherent: Lexical cohesion on a youtube channel. *Discourse, Context & Media*, 44, 100548. <https://doi.org/10.1016/j.dcm.2021.100548>
- Tazhenova, G., Mikhaylova, N., & Turgunbayeva, B. (2024). Digital media in informal learning activities. *Education and Information Technologies*. <https://doi.org/10.1007/s10639-024-12687-y>
- Weinberger, A., & Fischer, F. (2006). A framework to analyze argumentative knowledge construction in computer-supported collaborative learning. *Computers and Education*, 46. <https://doi.org/10.1016/j.compedu.2005.04.003>
- Wever, B. D., Schellens, T., Valcke, M., & Keer, H. V. (2006). Content analysis schemes to analyze transcripts of online asynchronous discussion groups: A review. *Computers and Education*, 46. <https://doi.org/10.1016/j.compedu.2005.04.005>
- Wise, A. F., Zhao, Y., & Hausknecht, S. N. (2013). Learning analytics for online discussions: A pedagogical model for intervention with embedded and extracted analytics. *Proceedings of the 3rd International Conference on Learning Analytics & Knowledge (LAK '13)*, 48–56. <https://doi.org/10.1145/2460296.2460308>
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., ... Rush, A. M. (2020). Transformers: State-of-the-art natural language processing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>
- Ye, D., & Pennisi, S. (2022). Analysing interactions in online discussions through social network analysis. *Journal of Computer Assisted Learning*, 38. <https://doi.org/10.1111/jcal.12648>

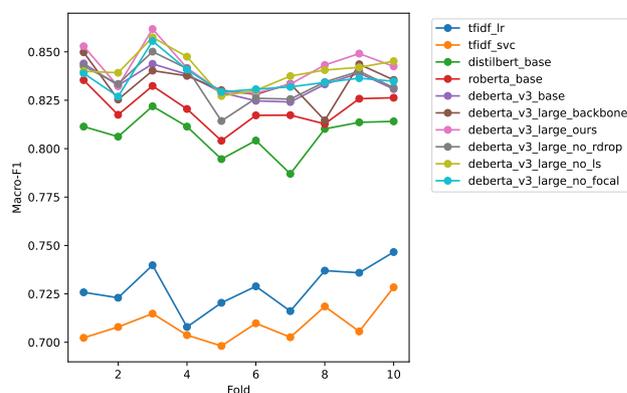
## A. Appendix A

**Table 10**  
Coding book for knowledge construction categories (Part 1).

Category	Indicator	Definition	Exclusion	Example
NonKC	Express emotion	Comments that express emotions or exaggerated reactions without reference to knowledge or reasoning.	If the emotion evaluates or comments on the scientific stance → Share / Explore.	"lol"
	Socialise or acknowledge others	Greetings, compliments, or support without content-related viewpoints.	If cognitive evaluation of content is present → Share / Explore.	"Love your channel—keep it up!"
	Produce irrelevant or incoherent content	Symbols, stretched letters, or random fragments lacking interpretable meaning.	If any meaningful knowledge-related content is included → Share / Explore.	"asdjkl..... !!"
	Ask unrelated questions to the learning content	Questions about the creator, platform, or unrelated topics.	If asking about content concepts → Share; if challenging → Explore.	"When's your TV show?"
Share	State an opinion or observation	Personal viewpoint or observation without conflict, reasoning, or synthesis.	If contradicts or refutes → Explore; if integrates → Integrate.	"Heat rises!"
	Express simple agreement	Agreement without justification or elaboration.	If reasons or evidence are added → Explore / Integrate.	"Totally agree!"
	Provide simple example or information	One additional fact or instance without argumentative intent.	If used to challenge or evaluate → Explore; if comparing sources → Integrate.	"The computer already has night-light settings."
	Ask for clarification	Questions seeking information but not challenging correctness.	If questioning accuracy or logic → Explore.	"Why is the balloon doing that and then it stayed like that?"
	Identify a problem or issue	Pointing out an issue or uncertainty about the topic without challenging or evaluating reasoning.	If directly challenges → Explore; if integrating solutions → Integrate.	"The image seems too blurry to tell what's actually happening."



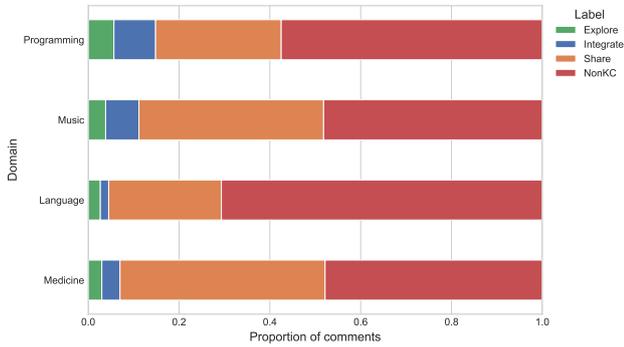
**Figure 3:** Distribution of 10-fold macro-F1 scores for all models. Each box shows the variability across folds, with individual fold scores overlaid as points. Transformer-based models consistently outperform the TF-IDF baselines, and the DeBERTa variants form a compact group of high-performing models.



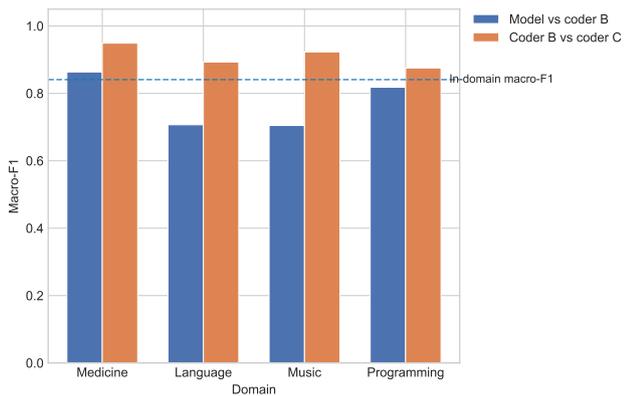
**Figure 4:** Fold-wise macro-F1 trajectories for all models. Each line corresponds to a model evaluated on the same 10 cross-validation folds. Transformer-based models dominate the TF-IDF baselines on nearly every fold, illustrating the repeated-measures structure of the analysis.

**Table 11**  
Coding book for knowledge construction categories (Part 2).

Category	Indicator	Definition	Exclusion	Example
Explore	Express disagreement	Disagreement or doubt without integrating multiple perspectives.	If synthesising or proposing a new framework → Integrate.	"You got frequency and wavelengths mixed up."
	Question or probe another statement Defend one's position with minimal support	Questions evaluating correctness, logic, or evidence. Defending a position with one or two reasons but without deeper comparison or integration.	If merely seeking information → Share. If multiple perspectives are weighed → Integrate.	"Isn't platysma just a de-generated version of this?" "Really? I heard it's Pauli exclusion, not Heisenberg uncertainty."
Integrate	Clarify concepts or terms	Refines or corrects conceptual understanding within a broader framework.	If only points out terminology → Share.	"This isn't about $F = ma$ . It's actually an example of inertia, where the force comes from the wood slowing down."
	Propose or synthesise integrative statements	Combines multiple viewpoints or pieces of evidence to propose a new explanation or compromise.	If just taking sides → Share.	"It's partly about buoyancy and partly about surface tension. The two together explain why it behaves differently in saltwater."
	Evaluate statements by comparing perspectives, experiences, or information sources	Compares sources, consequences, or perspectives to draw a judgment.	If only a single piece of evidence is used → Explore.	"Studies say blue light isn't the main cause, but many people still report less strain with filters, so the effect may be more behavioural than optical."
	Apply or reflect on new understanding	Shows changed understanding or connects content to broader implications.	If only noting that something was learned without explicit reflection → Share.	"I always thought going low-carb was the only way to manage blood sugar. But after seeing people compare research showing that increasing fibre can help just as much, I tried adding beans and whole grains instead of cutting carbs completely. What surprised me is that it actually feels easier to stick to, and my energy swings are way smaller now."

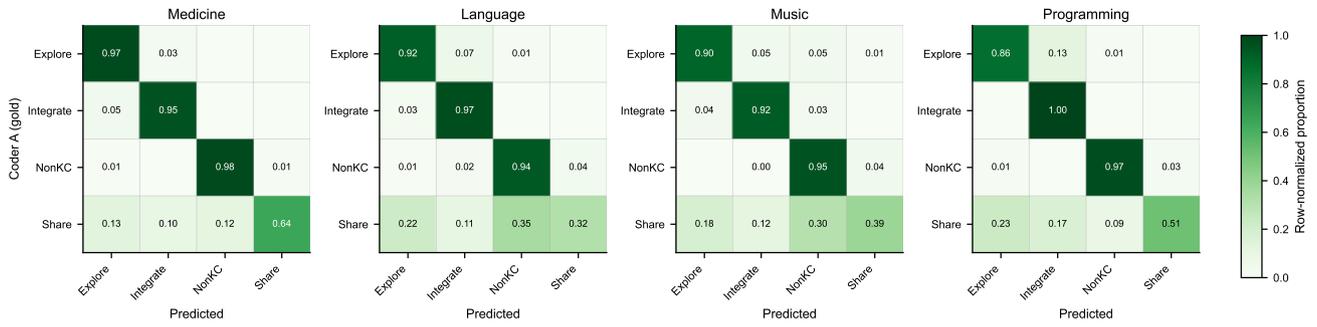


**Figure 5:** Natural knowledge construction category distribution by domain in the random subsets, based on coder B’s labels. Each horizontal bar shows the proportion of Explore, Integrate, Share, and NonKC comments within a domain.

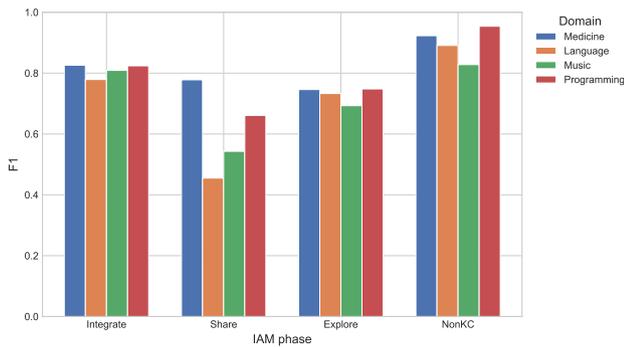


**Figure 6:** Cross-domain macro-F1 on the random subsets. Bars show macro-F1 for the classifier against coder B and for coder B against coder C in each domain. The dashed line indicates the in-domain macro-F1 of the classifier on the original dataset (Section 3.2).

## DeBERTa-KC: A Knowledge Construction Classifier



**Figure 7:** Confusion matrices of model-assigned categories across Medicine, Language, Music, and Programming, compared against coder B's labels (all samples: Random + Augment).



**Figure 8:** Per-category F1 scores by domain (all samples: Normal + Augment). Each group of bars corresponds to a knowledge construction category, with separate bars for Medicine, Language, Music, and Programming.