

Universal Design for Learning and Inclusive Classroom Participation

Madhubala Shrivastava

Assistant Professor

Sanskar City College of Education

Rajnandgaon, Chhattisgarh



<http://www.ujhmads.org/> || Vol. 2 No. 3 (2026): Special August Issue

Date of Submission: 02-07-2026

Date of Acceptance: 22-07-2026

Date of Publication: 17-08-2026

Abstract— Universal Design for Learning (UDL) has emerged as a major framework for reducing instructional barriers and accommodating learner variability in increasingly diverse classrooms. Despite growing adoption of accessible digital learning resources, relatively little research has examined whether UDL implementation can be connected to fine-grained patterns of classroom participation rather than conventional achievement or satisfaction outcomes alone. This study proposes an artificial-intelligence-assisted framework for investigating how multimodal representation, flexible engagement, and alternative modes of action and expression influence inclusive classroom participation. The central research gap concerns the absence of integrated models capable of examining behavioral, interactional, and accessibility indicators while accounting for differences in learners' preferred modes of participation. The proposed research conceptualizes participation as a multidimensional construct encompassing verbal contribution, digital interaction, collaborative involvement, task response, and sustained engagement rather than equating participation with speaking frequency.

Keywords— Universal Design for Learning, Inclusive Education, Classroom Participation, Learning Analytics, Explainable Artificial Intelligence, Multimodal Learning

Introduction

Contemporary classrooms are characterized by substantial variability in students' prior knowledge, linguistic backgrounds, cognitive strategies, communication preferences, accessibility requirements, motivation, confidence, and familiarity with digital technologies.

Instructional models designed around an assumed "average learner" therefore risk creating barriers even when the formal curriculum is available to every student. Inclusive education requires more than placing heterogeneous learners within a common physical or virtual classroom; it requires designing learning environments in which students have meaningful opportunities to access information, engage with learning activities, communicate understanding, collaborate with peers, and exercise agency.

Universal Design for Learning provides an influential framework for approaching this challenge proactively. Rather than waiting until a learner experiences difficulty and subsequently adding an accommodation, UDL encourages educators to anticipate learner variability during instructional design. The contemporary framework organizes this approach around multiple means of **engagement, representation, and action and expression**. CAST describes these dimensions as mechanisms for creating meaningful and challenging learning opportunities that are accessible to diverse learners.

The release of the **CAST UDL Guidelines 3.0 in July 2024** represents an important recent development. The revised framework places greater attention on learner agency and on barriers arising from bias and exclusionary educational structures. It also incorporates considerations related to belonging, identity, collaboration, accessible technologies, multiple forms of communication, and challenging exclusionary practices. These developments shift the discussion from accessibility understood primarily as access to instructional materials toward a broader conception in which learners should possess genuine opportunities to participate and exercise agency.

Classroom participation is particularly relevant to this expanded interpretation of inclusion. Participation influences the opportunities through which learners articulate ideas, receive feedback, negotiate meaning, collaborate, ask questions, demonstrate understanding, and develop academic confidence. Yet classroom participation is frequently measured through narrow indicators such as attendance, hand-raising, oral responses, or number of contributions. Such measurements can inadvertently privilege students whose communication preferences align with conventional classroom expectations.

A student who rarely speaks during whole-class discussion, for example, may contribute extensively through collaborative documents, anonymous polling, written discussion boards, peer-feedback activities, concept mapping, or small-group interaction. Conversely, frequent verbal participation does not necessarily demonstrate equitable or cognitively meaningful engagement. Consequently, inclusive participation should be understood as a **multichannel phenomenon** rather than as a synonym for classroom talk.

This distinction becomes increasingly important as digital technologies reshape instructional environments. Learning management systems, collaborative platforms, classroom response systems, captioned multimedia, speech-to-text tools, adaptive interfaces, and asynchronous discussion environments provide students with multiple pathways for interacting with academic material. Artificial intelligence has introduced an additional analytical layer because heterogeneous educational interaction data can potentially be examined to detect patterns that conventional classroom observation may overlook.

However, AI-supported educational analytics also introduce significant methodological and ethical risks. Models trained on conventional participation indicators could reproduce existing assumptions about what constitutes a "good" participant. Automated interpretation of facial expressions, disability status, emotion, attention, or cognitive capability is particularly problematic and should not be treated as a reliable basis for inclusive educational decision-making. An ethically defensible AI-supported UDL framework should therefore focus on observable learning interactions and environmental accessibility rather than attempting to infer sensitive characteristics.

Research Gap

Recent UDL research provides substantial conceptual and implementation-oriented evidence, but three analytical areas remain insufficiently integrated.

First, much of the literature evaluates UDL through achievement, accessibility, perceptions, course satisfaction, or broad engagement measures. A review by Cumming and Rose demonstrates the importance of UDL as an accessibility

framework in tertiary education while also illustrating the heterogeneous implementation and evidence base surrounding UDL research. Research examining students with intellectual disabilities has similarly reported promising academic, social, and behavioral outcomes while emphasizing the limited number and diversity of empirical studies.




Fig. 1: UDL Thinking

Source: <https://inclusive.tki.org.nz/guides/universal-design-for-learning/>

Second, participation is commonly represented by isolated indicators. Such operationalization can obscure learners who participate through less visible channels. There remains a need for a multidimensional participation construct that simultaneously considers verbal, written, digital, collaborative, task-oriented, and sustained forms of involvement.

Third, UDL research and educational machine-learning research have often developed along parallel trajectories. Learning analytics can detect complex interaction patterns, whereas UDL provides a theoretically grounded account of why multiple learning pathways should be designed. Yet comparatively little research has integrated these perspectives to determine whether specific UDL features are associated with more equitable distributions of classroom participation.

The present manuscript therefore addresses the following gap:

There is insufficient empirical research integrating UDL implementation indicators with multimodal learning analytics and explainable machine learning to evaluate whether inclusive instructional design broadens meaningful classroom participation across multiple

channels rather than merely increasing aggregate engagement.

This gap is particularly relevant following UDL Guidelines 3.0, whose emphasis on agency, belonging, identity, multiple communication modes, accessible technologies, and the removal of exclusionary barriers provides a stronger conceptual basis for investigating participation equity.

Research Objective

The principal objective of the proposed study is to develop and evaluate an **AI-assisted Inclusive Participation Analytics Framework (IPAF)** capable of examining relationships between UDL-oriented instructional design and multidimensional classroom participation.

Rather than predicting examination scores, the proposed model will estimate an **Inclusive Participation Index (IPI)** constructed from multiple observable interaction dimensions. The analytical architecture will investigate whether learners provided with flexible pathways for engagement, representation, and expression demonstrate broader and more balanced participation than learners experiencing comparatively rigid instructional environments.

A secondary objective is to use explainable machine-learning methods to determine which UDL-aligned instructional characteristics—such as response-mode flexibility, accessible multimedia, collaborative choice, asynchronous contribution opportunities, scaffolding, and alternative assessment formats—have the strongest associations with participation outcomes.

This orientation differs from conventional educational prediction studies because the intended question is not simply *which students are likely to participate?* Instead, it asks:

Which characteristics of the learning environment appear to create conditions in which participation becomes more accessible, diverse, and equitably distributed?

The distinction moves the analytical emphasis away from locating deficits within learners and toward identifying modifiable instructional barriers.

Literature Review

2.1 Evolution of Universal Design for Learning

UDL originated from the broader concept of universal design and developed into an educational framework concerned with proactively reducing barriers to learning. Its central premise is learner variability: students do not process information, sustain motivation, or demonstrate knowledge through identical pathways.

The contemporary UDL framework consequently emphasizes multiple means of engagement, representation, and action and

expression. CAST characterizes these dimensions as interconnected aspects of instructional design rather than isolated accommodations. Engagement concerns the reasons and conditions under which learners become involved; representation addresses alternative ways of perceiving and understanding information; and action and expression concern different ways learners interact with learning environments and communicate what they know.

UDL Guidelines 3.0 extend this orientation. Among the considerations incorporated into the framework are fostering belonging and community, optimizing accessible technologies, using multiple media for communication, supporting learner choice, addressing biases associated with modes of expression, and challenging exclusionary practices.

These changes are directly relevant to classroom participation. If learners are offered only one legitimate response channel—for example, immediate public speaking—participation opportunities may remain technically equal but functionally unequal. Providing several academically valid response pathways changes the architecture of participation itself.

2.2 UDL and Inclusive Educational Environments

The distinction between accommodation and universal instructional design is fundamental to understanding UDL. Reactive accommodation generally modifies an existing environment after a specific barrier has been identified. Universal design instead attempts to reduce predictable barriers during initial planning.

In higher education, Cumming and Rose reviewed UDL literature from an accessibility perspective and identified its potential to guide inclusive tertiary teaching while also highlighting variation in how UDL has been conceptualized and implemented. This heterogeneity remains methodologically important because merely labeling an intervention "UDL" does not demonstrate equivalent implementation fidelity.

A related systematic review focusing on students with intellectual disabilities identified interventions involving digital environments, e-books, shared stories, software, audio systems, and educational activities. The authors found evidence of positive educational effects but emphasized the limited empirical base and the need for broader research designs and larger samples.

The implication is that future UDL studies should measure the instructional features actually implemented rather than treating UDL as a binary variable. The present research therefore conceptualizes **UDL implementation intensity** as a multidimensional feature set instead of simply categorizing classrooms as "UDL" or "non-UDL."




Fig. 2: Pillars of UDL

Fig. 2: <https://www.structural-learning.com/post/understanding-universal-design-for-learning-udl>

2.3 From Accessibility to Participation Equity

Accessibility is a necessary condition for inclusion but is not sufficient evidence of participation. A digital resource can be technically accessible while the instructional activity surrounding it continues to restrict meaningful contribution.

For this reason, the proposed research distinguishes **participation quantity** from **participation equity**. Quantity concerns how much interaction occurs. Equity concerns whether legitimate opportunities for interaction are distributed across learners and communication modes.

Consider a classroom in which ten students produce 90% of verbal contributions. The overall number of classroom interactions might appear high, while participation remains concentrated among a small group. A UDL-oriented classroom may produce a different distribution by incorporating anonymous response systems, collaborative annotation, structured peer discussion, digital question submission, multimedia responses, and reflective writing.

The analytical challenge is therefore to measure not only whether participation increases but whether its **distribution becomes less concentrated and its modalities become more diverse**.

This manuscript introduces the concept of an Inclusive Participation Index to capture that distinction. The proposed index will integrate observable variables such as response frequency, response-mode diversity, collaborative contribution, task initiation, completion of optional interaction activities, peer-feedback involvement, and persistence across instructional sessions. The exact weighting and validation strategy are reserved for Part 2.

2.4 Multimodal Learning and Flexible Expression

Digital education has expanded the possible modes through which students can communicate academic understanding. Text, audio, video, diagrams, interactive quizzes, collaborative documents, presentations, simulations, and multimodal compositions can coexist within the same learning environment.

This development corresponds closely with UDL's emphasis on multiple means of action and expression. CAST Guidelines 3.0 explicitly address the use of multiple media for communication, multiple tools for construction and creativity, and flexible methods of interaction and response.

Importantly, flexibility does not imply lowering academic standards. The learning objective can remain stable while the pathway used to demonstrate competence varies. If the intended outcome concerns conceptual understanding rather than oral presentation ability, requiring every student to demonstrate that understanding exclusively through spontaneous speech may introduce an irrelevant participation barrier.

Accordingly, the proposed study treats modality choice as an instructional design variable rather than as a peripheral accessibility feature.

2.5 Artificial Intelligence and Learning Analytics for Inclusive Participation

Machine learning offers methods for identifying nonlinear relationships within complex educational datasets. Traditional regression remains valuable when relationships are approximately linear and theoretically specified, but classroom participation may depend on interactions among instructional structure, previous participation, response modality, task difficulty, collaboration opportunities, digital activity, and temporal patterns.

Tree-based ensemble models are particularly appropriate for the proposed study because they can accommodate mixed numerical and categorical variables while capturing nonlinear associations. Rather than using opaque deep-learning architectures simply for predictive performance, the planned framework prioritizes interpretability.

Explainable AI methods can subsequently estimate how specific instructional characteristics contribute to model predictions. In the context of UDL, this allows the research question to move beyond predictive accuracy toward instructional interpretation. For instance, an explainability analysis could determine whether availability of asynchronous responses, accessible multimedia, structured peer interaction, response-mode choice, or scaffolded task sequencing contributes most strongly to higher predicted participation.

Nevertheless, prediction must not be confused with causality. Feature importance indicates statistical contribution within a fitted model; it does not independently demonstrate that a particular UDL feature caused an improvement.

2.6 Responsible AI and the Risk of Algorithmic Exclusion

Applying AI to inclusive education requires particular caution because the objective of inclusion can be undermined by poorly designed analytics.

Models that infer attention from gaze, emotion from facial expression, or learning capacity from behavioral proxies may encode culturally and individually variable assumptions. Similarly, training a participation model primarily on spoken contributions would structurally disadvantage students who communicate through other channels.

The proposed architecture therefore follows three principles.

First, it analyzes **learning interactions rather than sensitive learner identities**. Disability diagnosis, medical information, facial recognition, and inferred emotional states are excluded from the predictive feature space.

Second, participation is represented through multiple channels. Written contributions, digital responses, collaboration, task interactions, and verbal participation are treated as legitimate components rather than positioning speech as the default standard.

Third, model interpretation is directed toward **environmental modification**. If a model identifies low participation probability, the intended response is not to label a student as disengaged. Instead, educators should examine whether additional response pathways, scaffolding, representation formats, collaboration structures, or accessibility options could reduce participation barriers.

This approach is consistent with the contemporary UDL emphasis on recognizing and challenging exclusionary structures rather than locating learning difficulties solely within individuals.

Proposed Methodology

Building on the research gap established in Part 1, this study proposes an **Inclusive Participation Analytics Framework (IPAF)** for evaluating how Universal Design for Learning (UDL) practices influence the breadth and equity of classroom participation. The framework deliberately shifts the analytical target from conventional academic-performance prediction toward the identification of instructional conditions associated with meaningful participation across multiple communication channels.

This orientation is consistent with CAST UDL Guidelines 3.0, which emphasize flexibility in response, communication,

interaction, accessible technologies, collaboration, belonging, and learner agency.

3.1 Research Design

A **quasi-experimental longitudinal design** is proposed across a 12-week instructional period. Participating classes are organized into two instructional conditions:

- **UDL-enhanced classrooms:** instructors systematically incorporate multiple means of engagement, representation, and action/expression.
- **Comparison classrooms:** instructors continue their existing instructional practices without being required to adopt the structured UDL intervention during the observation period.

Rather than defining UDL implementation as a simple binary variable, an instructor-level **UDL Implementation Fidelity Score (UIFS)** is calculated through structured observation and instructional-material audits. This distinction is important because implementation intensity can vary substantially even among classrooms formally described as UDL-oriented.

The intervention incorporates captioned and transcript-supported instructional media, adjustable digital learning resources, visual and textual representations, synchronous and asynchronous response channels, individual and collaborative tasks, optional response formats, structured peer interaction, and alternative methods for demonstrating understanding. CAST specifically recommends multiple media for communication and flexibility in methods of response and interaction.

The primary research question is:

Does greater implementation of UDL-aligned instructional design predict broader, more sustained, and more equitably distributed classroom participation when participation is measured across multiple interaction modalities?

Three working hypotheses guide the analysis:

H1: Students exposed to higher UDL implementation intensity will demonstrate significantly higher Inclusive Participation Index scores.

H2: UDL-enhanced classrooms will demonstrate lower concentration of participation across students than comparison classrooms.

H3: Response-mode flexibility and collaborative participation opportunities will emerge as stronger machine-learning predictors of inclusive participation than conventional lecture-oriented interaction variables.

Experimental Setup

4.1 Dataset Construction

Because no existing public dataset was identified that simultaneously contains classroom-level UDL implementation measures and the multimodal participation variables required by the proposed framework, the study uses a **purpose-built research dataset** rather than claiming use of an established dataset.

For the proposed experimental design, data are collected from **480 undergraduate learners distributed across 24 course sections** at participating higher-education institutions. Twelve sections implement the structured UDL intervention and twelve operate as comparison sections. Approximately 20 students are represented in each section.

The unit of longitudinal observation is the **student-week**, producing up to 5,760 observations before exclusion of incomplete records.

The dataset contains four major feature families:

Instructional-design features: UDL implementation score, number of representation formats, response-mode choice, availability of accessible materials, scaffolding intensity, asynchronous participation opportunities, and collaborative-learning opportunities.

Behavioral participation features: voluntary responses, prompted responses, questions submitted, peer-feedback events, group-task contributions, digital polls, discussion-board interactions, collaborative-document activity, and completion of optional learning activities.

Temporal features: response latency, participation consistency across weeks, inactivity intervals, session-to-session variation, and proportion of learning sessions containing at least one substantive contribution.

Contextual features: class size, instructional mode, task type, session duration, course discipline, week of instruction, and assessment proximity.

No facial-recognition data, biometric identifiers, medical information, disability diagnoses, inferred emotion labels, or automated estimates of cognitive ability are included. The framework is intended to evaluate modifiable characteristics of learning environments rather than classify students according to sensitive characteristics.

4.2 Inclusive Participation Index

A major methodological contribution is the construction of the **Inclusive Participation Index (IPI)**.

Participation is decomposed into five dimensions:

$$IPI_i = w_1V_i + w_2D_i + w_3C_i + w_4P_i + w_5S_i$$

where:

- (V_i) = verbal/structured response participation,
- (D_i) = digital interaction participation,
- (C_i) = collaborative contribution,
- (P_i) = participation-mode diversity,
- (S_i) = sustained participation across sessions.

Weights are initially estimated using confirmatory factor analysis and subsequently tested through sensitivity analysis rather than being assigned arbitrarily.

The resulting IPI is normalized to a **0–100 scale**, with higher values indicating broader and more sustained participation across legitimate learning channels.

A separate **Participation Concentration Index (PCI)** is calculated at classroom level using a normalized Herfindahl-type measure. A lower PCI indicates that classroom contributions are distributed across more students rather than being dominated by a small subset of highly visible participants.

This distinction permits a classroom to demonstrate high overall activity but poor participation equity—an outcome that conventional interaction counts could miss.

4.3 Data Preprocessing

All learner identifiers are replaced with randomized research codes before analysis. Raw event logs are aggregated into student-week observations to reduce unnecessary behavioral granularity.

Duplicate digital events generated by repeated system refreshes are removed. Extremely rapid repeated interactions that cannot reasonably represent separate learning contributions are consolidated according to predefined temporal thresholds.

Missing numerical values below a predefined tolerance are imputed using median values calculated exclusively from the training partition. Categorical missingness is represented through an explicit "unknown/not recorded" category where conceptually appropriate rather than silently assigning a dominant category.

Numerical predictors are standardized when required by the relevant algorithm. Highly skewed count variables are transformed using:

$$\begin{aligned} &[\\ &x' = \log(1+x) \\ &] \end{aligned}$$

Categorical variables are one-hot encoded for linear models, while native categorical handling is retained where supported by the selected boosting algorithm.

4.4 Model Architecture

The principal predictive model is **CatBoost regression**, selected because the proposed dataset contains heterogeneous numerical, categorical, ordinal, temporal, and interaction variables.

Three baseline models are included:

1. Elastic Net Regression;
2. Random Forest Regression;
3. HistGradientBoosting Regression.

CatBoost serves as the principal nonlinear model.

The model estimates:

$$[\hat{IPI}] = f(X_{UDL}, X_{interaction}, X_{temporal}, X_{context})$$

where $(\hat{IPI})$ represents predicted inclusive participation.

The purpose of the algorithm is not to assign fixed participation labels to students. Instead, it estimates the relationship between classroom design characteristics and participation outcomes after accounting for contextual and historical interaction features.

An additional classification experiment transforms the IPI into a predefined **low-participation-risk** outcome using thresholds determined on the training data. This secondary analysis evaluates whether the framework could identify learning environments in which additional participation pathways may be required.

4.5 Training Strategy

The dataset is divided at the learner level into approximately:

- 70% training data,
- 15% validation data,
- 15% held-out test data.

The final test partition remains untouched until model specification and hyperparameter selection are complete.

Grouped five-fold cross-validation is performed within the development data. Grouping by learner prevents repeated

weekly observations belonging to the same individual from crossing folds.

For CatBoost, hyperparameters considered include tree depth, learning rate, number of boosting iterations, L2 regularization strength, subsampling parameters, and early-stopping patience. Hyperparameter optimization is conducted exclusively on development data.

Early stopping is applied when validation performance fails to improve for a predetermined number of iterations.

To determine whether the complete multimodal feature architecture contributes meaningful information, three ablation experiments are also conducted:

$$[M_1 = \text{UDL design variables only}]$$

$$[M_2 = \text{UDL + conventional participation variables}]$$

$$[M_3 = \text{UDL + multimodal + temporal + contextual variables}]$$

The comparison is important because improved performance from (M_3) would provide evidence that inclusive participation cannot be adequately represented through speaking frequency or simple activity counts alone.

4.6 Evaluation Metrics

The regression task is evaluated using four complementary measures.

Mean Absolute Error (MAE):

$$\left[MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \right]$$

MAE provides an interpretable estimate of the average prediction error in IPI units.

Root Mean Squared Error (RMSE):

$$\left[RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \right]$$

RMSE penalizes larger errors more heavily.

Coefficient of Determination (R^2):

$$\left[R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2} \right]$$

This indicates the proportion of variation in participation captured by the model on unseen data.

Spearman Rank Correlation (ρ) is additionally calculated to determine whether predicted IPI scores preserve

the relative ordering of participation levels when strict linearity cannot be assumed.

For the secondary risk-classification analysis, **balanced accuracy, macro-F1 and area under the precision-recall curve (PR-AUC)** are used. PR-AUC is included because a low-participation-risk category may be less frequent than the non-risk category.

Results and Discussion

Important methodological note: the numerical findings below constitute a **proposed/simulated experimental scenario for manuscript development**, not results from an actually conducted human-participant study. They must be replaced by empirical outputs before submission as a completed experimental study.


The simulated evaluation indicates that the nonlinear models capture inclusive-participation patterns more effectively than the linear baseline. CatBoost produces the strongest overall test performance, suggesting that the relationship between UDL implementation and classroom participation is likely to involve nonlinearities and interactions rather than a simple additive structure.

Table 1. Simulated Comparative Performance of Inclusive Participation Models

Analytical Model	MAE ↓	RMS E ↓	Test (R^2) ↑	Spearman (ρ) ↑
Elastic Net	8.91	11.74	0.57	0.72
Random Forest	7.16	9.48	0.71	0.81
HistGradientBoosting	6.74	8.96	0.75	0.84
CatBoost	6.21	8.37	0.79	0.87

Model performance for inclusive participation prediction

Simulated test-set R^2 values for predicting the Inclusive Participation Index (IPI).



The simulated (R^2) of 0.79 for CatBoost would indicate that the integrated feature set accounts for a substantial proportion of variation in the constructed participation measure. More importantly, the difference between the linear and nonlinear models suggests that participation may depend on combinations of instructional characteristics.

For example, merely providing digital discussion tools may have limited value when students have little autonomy over response mode. Conversely, response-mode choice may become more useful when accessible representation and collaborative structures are simultaneously available.

This interpretation corresponds with the architecture of UDL itself. CAST does not present a single mode of engagement or expression as optimal for every learner; the framework instead emphasizes flexible options for engagement, representation, response, interaction, and communication.

5.1 Effect of UDL Implementation

In the simulated intervention analysis, UDL-enhanced classrooms demonstrate higher average IPI values and lower participation concentration than comparison classrooms.

The most important conceptual result is not simply an increase in total interaction. Instead, simulated UDL classrooms show a greater proportion of learners participating through at least two distinct channels.

This supports the proposed distinction between **activity amplification** and **participation inclusion**.

An instructional technology that causes already-active students to submit more responses increases activity but may not improve inclusion. In contrast, an instructional design that enables previously less-visible learners to contribute through writing, polling, collaborative annotation, asynchronous discussion, or alternative response formats broadens participation.

CAST's current guidance explicitly emphasizes honoring different methods of response and providing multiple modalities for expression and communication, which supports this multidimensional operationalization of participation.

5.2 Explainable AI Findings

In the simulated SHAP analysis, the strongest positive instructional predictors are:

response-mode flexibility, followed by **structured collaborative opportunities**, **availability of asynchronous contribution**, **multimodal representation**, and **accessible digital-resource availability**.

Response-mode flexibility emerging as the strongest predictor is theoretically meaningful. Students differ in how

effectively they communicate understanding, and a participation system dependent almost entirely on immediate oral responses can unnecessarily restrict the observable participation of students who demonstrate knowledge more effectively through other modes.

CAST Guidelines 3.0 similarly state that different forms of action and expression should be honored and that no single means of expression is optimal for every learner.

The explainability results should nevertheless be interpreted as associations rather than causal effects. SHAP values explain the behavior of a fitted predictive model; they do not independently establish that changing a particular instructional feature will cause a corresponding change in participation.

5.3 Ablation Analysis

The simulated ablation experiment provides additional support for the multidimensional framework.

The UDL-only model produces moderate predictive performance. Adding conventional classroom interaction indicators substantially improves prediction. The complete multimodal model, incorporating digital, collaborative, temporal, and contextual features, produces the strongest performance.

This result would imply that inclusive classroom participation is poorly represented by a single observable behavior.

The finding also has methodological significance. Educational analytics systems that define engagement primarily through login counts, attendance, speaking frequency, or assignment completion risk overlooking meaningful but less visible forms of participation.

5.4 Educational Implications

The proposed framework has three practical implications.

First, instructors could use aggregated analytics to determine whether classroom participation is concentrated among a relatively small proportion of students.

Second, educators could examine **participation pathways** rather than interpreting low verbal activity automatically as disengagement.

Third, explainability outputs could support instructional redesign. If asynchronous opportunities consistently correspond with broader participation, instructors might expand structured discussion boards or pre-class question submission. If collaborative scaffolding appears important, group activities could be redesigned to distribute meaningful roles rather than merely placing students into groups.

The model should therefore function as a **design-support system**, not an automated decision-maker.

Future Scope

Future research should validate the framework across secondary schools, universities, blended programs, fully online environments, and different disciplinary contexts.

A particularly important direction would involve **privacy-preserving federated learning**, through which participating institutions could contribute to model development without centrally pooling identifiable student-level data.

Future studies could also investigate causal effects using cluster-randomized designs in which classrooms are assigned to different UDL implementation strategies.

Longitudinal research should examine whether expanded participation pathways gradually influence learner agency, persistence, belonging, and academic self-regulation.

Another direction concerns adaptive learning environments. Rather than predicting which learner has a "deficit," an AI system could recommend changes to the **learning environment**—for example, adding an asynchronous response channel, providing another representation format, increasing scaffolding, or offering an alternative mode of expression.

Such a direction is more consistent with the UDL principle of removing environmental barriers than designing algorithms that attempt to normalize learner behavior. CAST's 2024 update explicitly strengthened attention to barriers rooted in bias and systems of exclusion.

Conclusion

This study proposes a distinct intersection between **Universal Design for Learning, inclusive education, multimodal learning analytics, and explainable machine learning**. Rather than treating classroom participation as speaking frequency or aggregate digital activity, the proposed Inclusive Participation Analytics Framework conceptualizes participation as a multidimensional phenomenon involving verbal, digital, collaborative, modal, and longitudinal engagement.

The framework further changes the conventional objective of educational prediction. Its purpose is not to determine which students conform most closely to a predetermined model of classroom engagement. Instead, it seeks to identify which **instructional conditions reduce barriers and expand legitimate pathways for participation**.

The proposed CatBoost-based architecture combines UDL implementation indicators with multimodal interaction and temporal features, while SHAP-based interpretation provides

a mechanism for identifying influential instructional characteristics. The simulated results suggest that response flexibility, collaborative opportunities, asynchronous interaction, multimodal representation, and accessible resources could provide substantial explanatory value when inclusive participation is modeled as an educational outcome.

The manuscript also establishes a critical boundary for responsible AI: classroom analytics should not transform learner variability into algorithmic deficiency. Systems designed for inclusive education should direct analytical attention toward modifiable environmental barriers and should preserve multiple legitimate forms of engagement and expression.

This approach is particularly compatible with UDL Guidelines 3.0, which emphasize learner agency, flexible response and communication, accessible technologies, belonging, collaboration, and the reduction of exclusionary barriers.

References

1. CAST. (2024). *Universal Design for Learning Guidelines version 3.0*. CAST. UDL Guidelines 3.0 was released on July 30, 2024, and explicitly expands attention to learner agency, bias, exclusion, belonging, accessible technologies, communication, and multiple means of action and expression.
2. Cumming, T. M., & Rose, M. C. (2022). Exploring universal design for learning as an accessibility tool in higher education: A review of the current literature. *The Australian Educational Researcher*, 49, 1025–1043. <https://doi.org/10.1007/s13384-021-00471-7>
3. AlRawi, J. M., & AlKahtani, M. A. (2022). Universal design for learning for educating students with intellectual disabilities: A systematic review. *International Journal of Developmental Disabilities*, 68(6), 800–808. <https://doi.org/10.1080/20473869.2021.1900505>
4. Meyer, A., Rose, D. H., & Gordon, D. (2014). *Universal Design for Learning: Theory and Practice*. CAST Professional Publishing.
5. CAST. (2018). *Universal Design for Learning Guidelines version 2.2*. Wakefield, MA: CAST.
6. Bergdahl, N., Bond, M., Sjöberg, J., Dougherty, M., & Oxley, E. (2024). Unpacking student engagement in higher education learning analytics: A systematic review. *International Journal of Educational Technology in Higher Education*, 21, 63. <https://doi.org/10.1186/s41239-024-00493-y>
7. Bond, M., Viberg, O., & Bergdahl, N. (2023). The current state of using learning analytics to measure and support K-12 student engagement: A scoping review. In *Proceedings of the 13th International Learning Analytics and Knowledge Conference (LAK23)* (pp. 240–249). ACM. <https://doi.org/10.1145/3576050.3576085>
8. Dorogush, A. V., Ershov, V., & Gulin, A. (2018). CatBoost: Gradient boosting with categorical features support. *arXiv preprint arXiv:1810.11363*.
9. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, 4765–4774.
10. Bergdahl, N., Bond, M., Sjöberg, J., Dougherty, M., & Oxley, E. (2024). Unpacking student engagement in higher education learning analytics: A systematic review. *International Journal of Educational Technology in Higher Education*, 21, 63. <https://doi.org/10.1186/s41239-024-00493-y>
11. Bond, M., Viberg, O., & Bergdahl, N. (2023). The current state of using learning analytics to measure and support K-12 student engagement: A scoping review. *Proceedings of LAK23*, 240–249. <https://doi.org/10.1145/3576050.3576085>
12. Synthesis of student engagement with digital technologies: A systematic review of the literature. (2021). *Education and Information Technologies*. The review discusses how digital learning environments and LMSs provide multiple synchronous and asynchronous interaction opportunities and generate data suitable for learning analytics.
13. *A Comprehensive Review of Multimodal Analysis in Education*. (2025). *Applied Sciences*, 15(11), 5896. The review examines multimodal learning analytics using diverse educational data sources and highlights interpretability, ethics, synchronization, and scalability challenges.