ISSN 2409-2665
Journal of Logistics, Informatics and Service Science
Vol. 12 (2025) No. 3, pp. 141-162
DOI:10.33168/JLISS.2025.0310

Predicting Faculty Engagement in Communities of Practice: A Neural Network Approach with Synthetic Data Augmentation

Hamed Hilal Nasser Al Yahmadi, Yousuf Nasser Al Husaini*, Sami Khatar Al Mazruii, Walid Aboraya

Faculty of Education Studies, Arab Open University, Muscat, Oman

hamed.y@aou.edu.om, yousufnaser@aou.edu.om (Corresponding Author), sami.m@aou.edu.om, walid.aboraya@aou.edu.om

Abstract. Communities of Practice (CoPs) offer valuable faculty development platforms in higher education, yet their evaluation is hindered by small sample sizes and predominantly descriptive analytics. This study addresses this methodological gap by developing a deep learning framework enhanced by data augmentation techniques to predict faculty engagement in CoPs. We collected survey data measuring six CoP dimensions from 29 academic staff at Arab Open University-Oman and expanded this limited dataset using Synthetic Minority Oversampling Technique (SMOTE), categorical permutation, and Gaussian noise injection. The resulting neural network model achieved 89% accuracy and an F1-score of 0.89, significantly outperforming both baseline logistic regression (69%) and non-augmented deep learning (82%) approaches. Feature importance analysis identified shared leadership and peer collaboration as the strongest predictors of engagement, while t-SNE visualizations revealed distinct behavioral clusters among faculty. This study contributes to educational research by demonstrating that (1) deep learning with appropriate data augmentation can effectively model faculty engagement patterns even with limited initial data, and (2) CoP theoretical constructs can be computationally validated and operationalized, enabling more targeted faculty development interventions. The framework provides institutions with a scalable method for evidence-based decision-making in professional development planning, helping identify disengaged faculty and prioritize high-impact support strategies.

Keywords: Faculty Development, Communities of Practice, Deep Learning, Data Augmentation, Educational Analytics, Predictive Modeling, Neural Networks, SMOTE, Higher Education, Institutional Planning

1. Introduction

College staff development remains vital because it sustains educational excellence, enhances innovation, and supports perpetual learning in college settings (Hénard & Roseveare, 2012). Universities directed toward data-driven environments make it essential to grasp faculty needs alongside evaluating professional development practices that include CoPs in complex ways. Educational Communities of Practice consist of problem-solving groups that support educators in pursuing joint learning and mutual understanding development. These communities support the development of leadership abilities and innovative thought processes, which extend to the betterment of academic leadership. Traditional faculty development assessment relies on surveys, interviews, and performance reports. While these methods provide valuable qualitative insights, they typically yield small, imbalanced datasets that limit statistical analysis and generalizability. These valuable data collection methods lead to datasets that are restricted in depth and unable to generalize findings because they create imbalanced and small datasets (Guskey, 2000). Subjective characteristics of this data set obstacle to establishing engagement prediction patterns through time metrics. Artificial intelligence has entered a new era through deep learning, which provides encouraging solutions for various applications. Deep learning models use their power to analyze extensive, ill-structured, complicated data, which reveals hidden patterns that traditional analysis methods cannot detect (LeCun et al., 2015). These models achieve successful outcomes when applied to educational datasets, as they identify risks among students and create performance predictions and behavioral classes (Alhazmi & Sheneamer, 2023; Hernández-Blanco et al., 2019). Research about deep learning applications to evaluate teaching development strategies, particularly in the context of CoP frameworks, remains underexplored. Using deep learning properly on educational survey data requires large datasets because they determine the success of analysis. Small faculty development datasets cause performance issues for models and prevent their ability to generalize useful information. Data augmentation proves essential for this operation. The concept emerged from image processing, but data augmentation now serves text and tabular datasets by creating synthetic data for training expansion, model accuracy improvement, and reduced overfitting (Rizos et al., 2019; Shorten & Khoshgoftaar, 2019). This study aims to create a modern computational method for assessing faculty engagement within CoPs using deep learning and data augmentation techniques. Through this integration, researchers gain more profound perceptions of data patterns and hidden relationships, which provides institutions with valuable data-driven decisions for professional learning programs. These frameworks provide a union between educational conventions and advanced AI techniques to create better flexible analysis methodologies for faculty development. Despite increased investment in faculty development initiatives globally, their effectiveness remains inconsistent. According to an OECD survey (2021), over 52% of university faculty report inadequate support in improving teaching and engagement practices. Similarly, UNESCO (2022) highlights that nearly half of professional development programs in higher education lack follow-up or measurable outcomes, resulting in minimal impact on long-term instructional quality. These findings underline the urgency of designing evidence-based, scalable tools to support faculty growth, particularly in distributed, resource-constrained institutions.

Quality teaching depends on faculty development since it promotes innovation and enables continuous university learning. Higher education entities transition toward data-secured operational frameworks and thus demand heightened knowledge of faculty requirements and evaluating practices such as CoPs, which have become intricate and pressing. Educators come together in CoPs to jointly learn while building knowledge through collaboration. The establishment of CoPs enables academic communities to develop leadership as well as reflective practice and innovative approaches to teaching. Previous information acquisition regarding faculty development relied on survey responses, performance reports, and face-to-face interviews. These methods produce limited, small-sized, unbalanced datasets, which restrict analytical depth, along with the generalizability of obtained results (Guskey, 2000; Hsieh, 2010). Such data tends to have subjective elements that impede researchers' ability to find patterns and engage in time-related predictions. Deep learning innovations that originate from artificial intelligence research provide promising new options. The analysis and diverse capacity of deep learning models allow the extraction of obscured data patterns from big, unstructured, complicated datasets (Goodfellow et al., 2016). Educational data outcomes from these predictive and behavioral classification models demonstrate success in performance assessment risk identification and learning type classification (Abubakaria et al., 2020; Li & Liu, 2021). Deep learning solutions for evaluating faculty development practices within CoP frameworks have received limited investigation in academic research. An issue with deep learning applications to educational survey data comes from the limited sizes of available datasets. The size of available faculty development datasets creates modeling performance issues that restrict their generalization potential. The effective implementation of data augmentation becomes necessary at this point. The image processing technique named data augmentation now serves text and tabular data through expansion techniques, which assist in producing larger training sets along with minimizing overfitting while enhancing accuracy predictions (Shorten & Khoshgoftaar, 2019; Ying, 2019). The current work develops a new method to study faculty involvement with CoPs using complex artificial intelligence approaches, which include deep learning and data augmentation. The integrated process allows researchers to discover significant and concealed data patterns that help institutions make data-driven choices for professional learning development at scale. Such research connects educational science with modern AI technologies, which leads to advanced analytical systems that can optimize faculty development programs.

Wide acknowledgement exists for CoPs as providers of professional learning platforms in higher education, yet measuring their effectiveness through standardized data analytics systems remains challenging. Survey methods and qualitative response platforms continue to be the main tools for measuring faculty engagement since most institutions follow these traditional assessment methods (Day & Sachs, 2005). These useful assessment methods produce tiny and structurally skewed data sets that restrict scientists from executing detailed and predictive research (Arthur, 2016). The training process for advanced machine learning models requires continuous data that should include numerical variables because categorical and ordinal responses typically present themselves in limited training scenarios. The success of deep learning in healthcare, along with finance and personalized education, demands abundant datasets containing diverse information for optimal operation (Bengio et al., 2017). The use of these techniques on small educational survey data leads to overfitting that reduces generalization potential (Feng et al., 2019). The problem of scarce data requires innovative solutions, and researchers now apply data augmentation methods in natural language processing and tabular learning domains (Shorten & Khoshgoftaar, 2019). The methods produce synthetic growth of datasets through a process that generates more diverse data, which enables deep models to uncover more resilient patterns (Rizos et al., 2019). Data augmentation techniques have been underutilized in faculty development research, with limited empirical testing of their effectiveness in CoP evaluation contexts. Traditional research approaches for faculty development cannot detect complex nonlinear patterns between variables such as leadership commitment, institutional backing, and peer relationships (Hsieh, 2010). Universities that lack predictive modelling tools operate reactively when serving faculty, leading them to miss crucial opportunities to identify low-performing and high-performing staff members. According to this research, a framework combining data augmentation methods with deep learning techniques would analyze CoP-related survey data, filling significant knowledge gaps. The framework aims to achieve deep analysis through enhanced predictions that generate usable information for academic development leaders to make strategic choices. While prior research has explored student performance prediction using machine learning, there is limited empirical work on modeling faculty engagement, especially in the context of Communities of Practice (CoPs). Existing studies either rely on large institutional datasets or lack predictive depth. Moreover, the intersection of AI-driven modeling and CoP engagement theory remains largely unexplored, particularly under constraints of small, imbalanced, and survey-based datasets. This study addresses this gap through an interpretable deep learning framework informed by CoP theory.

This research study has the following specific objectives:

- To preprocess and structure CoP-related survey data collected from higher education faculty, ensuring it is suitable for computational modeling.
- To apply data augmentation techniques—such as synthetic minority oversampling, categorical feature permutation, and noise injection—to expand and balance the original dataset.
- To design and train a deep learning model, particularly a multi-layer neural network, capable of identifying patterns in faculty responses and predicting levels of engagement or institutional readiness for CoP integration.
- To interpret the model outputs in the context of faculty development, offering practical insights into institutional strengths, potential gaps, and targeted policy actions.

By achieving these objectives, the study aims to demonstrate how AI-driven models can complement educational theory and enhance institutional decision-making in faculty development planning. This paper contributes to academic research and applied artificial intelligence by proposing an innovative framework for evaluating faculty development initiatives, particularly CoPs, using advanced deep learning and data augmentation techniques.

This paper contributes to both educational theory and applied artificial intelligence in the following ways:

- It presents a novel methodology that combines three data augmentation strategies (SMOTE, categorical permutation, and Gaussian noise injection) with a deep learning model tailored for low-resource educational contexts.
- The approach enables predictive modeling of faculty engagement, offering interpretable outputs that align with Wenger's CoP dimensions, such as mutual engagement and shared repertoire.
- From a practical perspective, the model supports institutional decision-making by identifying latent engagement patterns and surfacing key development needs, particularly in under-resourced or distributed universities.

These contributions demonstrate how AI can enhance precision in professional development planning while remaining grounded in educational theory and ethical implementation.

The remainder of this paper is organized as follows. Section 2 reviews related work on CoPs, faculty development, and the application of deep learning and data augmentation in educational analytics. Section 3 outlines the methodology, including data collection, augmentation techniques, model architecture, and evaluation strategies. Section 4 presents the results, including performance metrics, visualizations, and learned pattern analysis. Section 5 discusses the findings in the context of faculty development theory and institutional practice. Finally, Section 6 concludes the paper with key insights, practical recommendations, limitations, and directions for future research.

2. Related Work

2.1 Communities of Practice and Faculty Development

CoPs, introduced by Lave and Wenger, describe groups of individuals who learn together through shared practices and experiences (Lave & Wenger, 1991; Wenger, 1999). In higher education, CoPs have been adopted to foster professional learning, innovation in teaching, and faculty collaboration (Cox & Richlin, 2004). They support the development of shared values, reflective practice, and collective inquiry (Day & Sachs, 2005). Arthur emphasizes that CoPs are especially beneficial in academic environments, where learning is often siloed and self-directed (Arthur, 2016). CoPs have also been associated with improved instructional design, peer mentoring, and knowledge-sharing (Prenger et al., 2019). In the Gulf region, Alyahmadi and Al-Sammakhi highlighted the potential of CoPs to align institutional goals with faculty development needs, though implementation remains inconsistent (Al-Yahmadi & Al-Shammakhi, 2021). Despite their value, most CoP evaluations rely on interviews, self-report surveys, and anecdotal reflections (Day, 2018; Guskey, 2000). While necessary, these tools lack scalability and often fail to reveal deeper behavioral or engagement patterns. There is growing interest in applying computational methods to improve the evaluation of such learning communities, but few studies have done so within the context of faculty development.

2.2 Faculty Development and Educational Data Analytics

Faculty development initiatives typically include workshops, mentoring, performance reviews, and peer observation. Studies show that sustained, peer-led approaches are more effective than top-down, episodic training (Chalmers & Keown, 2006; Farrell, 2015) However, most universities still use descriptive analytics to monitor development activities, limiting their ability to generate actionable insights (Macfadyen & Dawson, 2010). The use of analytics in education has been primarily focused on students, monitoring engagement, predicting dropout, and personalizing learning pathways (Ferguson, 2012; Yau et al., 2018). Faculty data applications are often limited due to smaller datasets and privacy concerns. Nonetheless, recent work suggests that predictive analytics can help institutions identify which faculty members need additional support and which development practices are most impactful (Leitner et al., 2017; Siemens, 2012). Calls have been made for integrating AI techniques in faculty development evaluation to move from descriptive to predictive and prescriptive insights (Okewu et al., 2021). However, this shift requires robust data preparation and model design, especially in sparse or imbalanced data contexts.

2.3 Deep Learning and Data Augmentation

Deep learning has shown strong performance in education-related applications, including text classification, student performance prediction, and learning analytics (Liu et al., 2021). Models such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformers have been applied to learning management systems, feedback systems, and student essays (Rusk, 2016; Shorten & Khoshgoftaar, 2019). A wide range of prediction methods for academic performance have been reviewed, covering both traditional and AI-based approaches (Al Husaini & Shukor, 2022). A significant barrier to using deep learning in education is the limited availability of high-quality, labeled data. This is particularly true in faculty development research. To overcome this, data augmentation has emerged as a strategy to expand datasets while synthetically preserving data structure and semantics. Originally popular in image and NLP domains, these techniques are now adapted for structured survey and tabular data (Marivate & Sefara, 2020; Mujahid et al., 2024). SMOTE (Synthetic Minority Oversampling Technique), random noise injection, categorical permutation, and GAN-based synthesis have been applied to augment small-scale educational data (Chawla et al., 2002; Fang et al., 2022). Kobayashi proposed a novel framework for augmenting tabular and textual data in learning systems, significantly improving model generalization (Seltzer, 1991). Deep learning models have also been successfully applied in virtual learning environments to predict academic performance, as demonstrated by Al Husaini and Shukor (Al Husaini & Shukor, 2024). Despite these advances, deep learning and data augmentation for analyzing faculty CoP engagement are virtually unexplored. This paper builds on foundational work in educational AI and offers a novel integration of these methods in the context of academic professional development.

2.4 Predictive Modeling of Faculty Engagement

Predictive modeling in educational contexts has historically centered around student outcomes, such as academic performance, dropout prediction, and behavioral analytics. While these applications have yielded valuable insights, comparatively fewer studies have focused on modeling faculty engagement, particularly within professional development environments such as Communities of Practice (CoPs). Faculty engagement is a multi-dimensional construct influenced by institutional support, leadership, collaboration, and reflective teaching practices. These aspects are typically captured through surveys and feedback instruments—datasets that are often small, imbalanced, and difficult to generalize. Unlike student datasets, which may include log files and longitudinal academic records, faculty-related data lack behavioral granularity and are often self-reported. This poses challenges for building robust, interpretable predictive models. Few existing works have attempted to operationalize or quantify faculty engagement using artificial intelligence. For example, prior research by Al Husaini and Shukor (Al Husaini & Shukor, 2024) successfully applied deep learning techniques to predict student academic performance in virtual learning environments. However, these studies relied on larger behavioral datasets and did not address professional development or faculty engagement metrics. Their methodologies also do not tackle data imbalance or survey-specific noise, which are common in faculty development research. Recent advances in data augmentation for tabular data offer promising solutions to this gap. Studies such as (Yadav et al., 2025) introduced conditional GANs (CTGAN) for synthesizing structured data in low-sample settings, while (Chawla et al., 2002) demonstrated the utility of SMOTE and hybrid techniques for educational surveys. These methods enhance the model's ability to generalize without overfitting on sparse or skewed features. Despite these advancements, little work has been done to combine deep learning with targeted augmentation techniques specifically for faculty engagement datasets. Our study addresses this gap by designing a framework that integrates SMOTE, categorical permutation, and Gaussian noise injection to enrich a limited CoP survey dataset. This enables accurate classification of engagement levels while preserving interpretability through feature importance and visualization. By modeling faculty engagement with a specialized neural architecture trained on synthetically expanded data, this study contributes to the emerging field of precision education analytics for academic staff—an area where scalable, AI-driven decision tools are increasingly needed.

3. Methodology

This study adopts an applied research methodology integrating data augmentation techniques and deep learning models to analyze faculty development data and precise responses related to CoPs. The approach is based on a quantitative, model-driven pipeline supported by systematic data preparation and evaluation stages. The research methodology can be studied through Fig. 1, which displays all the sequential steps of this investigation. The method starts with collecting survey data about key dimensions of CoP and then moves to encoding and normalization preprocessing steps. The dataset gets expanded through augmentation techniques that feature SMOTE and noise injection. The training procedure utilizes a fully connected deep neural network to perform the task of engagement level classification. Evaluation takes place after the model implementation through visualization of data embeddings and feature importance assessments that lead to actionable insights for faculty development.

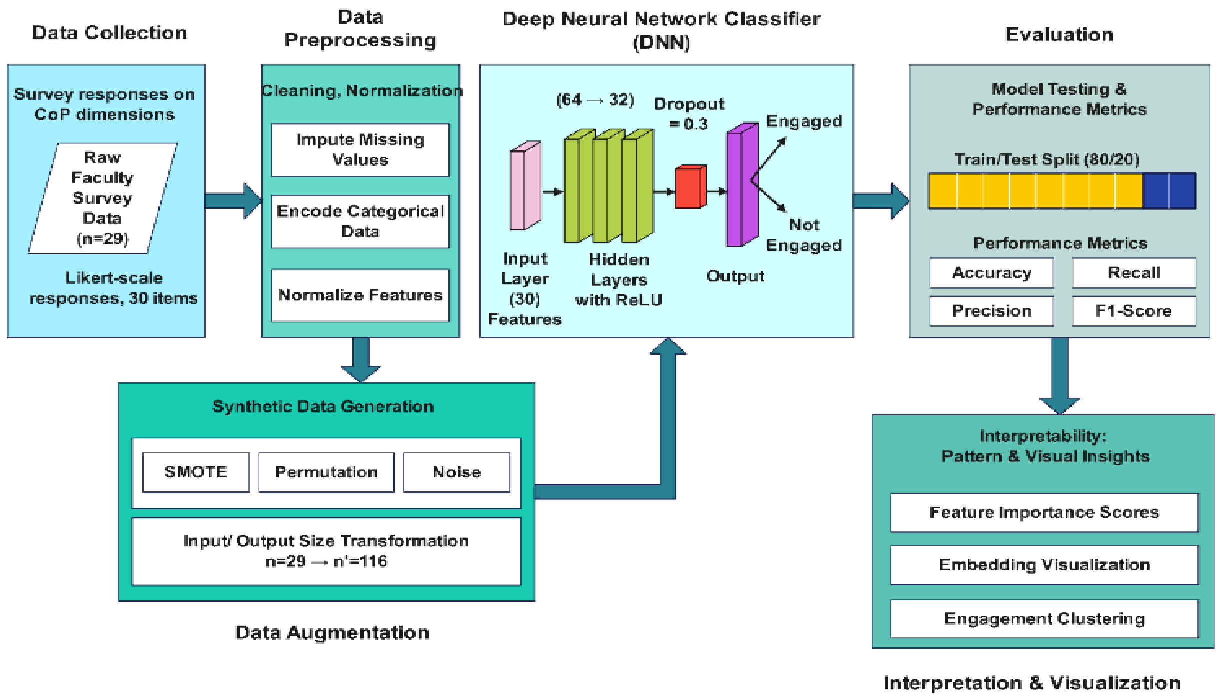

Fig. 1 Proposed Methodology for Predicting Faculty Engagement in CoPs

3.1 Data Collection and Preprocessing

For this research, data were collected through an online structured survey aimed at academic staff members, policy decision-makers, and students attending the Arab Open University—Oman. The research instrument measured CoP dimensions through six main characteristics: shared leadership, shared values, collective learning, personal practice and relationships, and structural support. Each dimension included multiple Likert-scale items (1 = Strongly Disagree to 5 = Strongly Agree). A total of 29 participants responded to the survey. The raw data included 30 categorical items, each corresponding to a statement measuring perception or experience related to CoPs. Responses were transformed into numerical form using ordinal mapping:

$$\text{Strongly Disagree} = 1, \text{Disagree} = 2, \text{Neutral} = 3, \text{Agree} = 4, \text{Strongly Agree} = 5$$

Let the dataset be represented as a matrix $X \in R^{n \times d}$, where: $n = 29$ is the number of participants, $d = 30$ is the number of survey items (features). Each entry x_{ij} in the matrix represents the numerical response of the participant i to item j, where $x_{ij} \in \{1,2,3,4,5\}$.

Missing or inconsistent responses (e.g., blanks) were handled using mean imputation for each item. Let μ_j be the mean of the non-missing values in column j, then for any missing value $x_{ij} = NaN$, it was replaced with:

$$x_{ij} = \mu_j \tag{1}$$

To prepare the data for deep learning, the input features were normalized using Min-Max scaling to bring all values into the [0, 1] range:

$$x'_{ij} = \frac{x_{ij} - min\,(x_j)}{max\,(x_j) - min\,(x_j)} \tag{2}$$

This ensured that no single variable dominated the learning process due to scale differences. Table 1 summarizes the class distribution before and after applying data augmentation. The original dataset exhibited a class imbalance, with 13 instances of Low Engagement and 16 of High Engagement. Through SMOTE and permutation-based augmentation, the dataset was expanded and balanced to enhance model training.

Table 1. Class Distribution Before and After Augmentation

Class Label	Original Count	After Augmentation
Low Engagement	13	58
High Engagement	16	58
Total	29	116

3.2 Data Augmentation Strategies

Data augmentation was employed to address two core challenges in the dataset: small sample size and class imbalance. Although most data augmentation techniques have been developed for unstructured data such as images or text, recent advances have enabled their adaptation to tabular and categorical educational data. This study applied three augmentation techniques: SMOTE, random categorical permutation, and Gaussian noise injection. Each method was intended to expand the dataset and improve model generalization.

SMOTE was applied to synthetically generate new samples for underrepresented CoP engagement patterns in the survey data. Given a minority class C_m, new samples were generated by linear interpolation between an existing instance $x_i \in C_m$ and one of its k-nearest neighbors $x_k \in C_m$, defined as:

$$\tilde{x} = x_i + \lambda \cdot (x_k - x_i), \lambda \sim U(0,1) \tag{3}$$

Where, λ is a random scalar from the uniform distribution, $\tilde{x}$ is the newly generated synthetic instance. The number of nearest neighbors (k) in SMOTE was set to 5, following standard practices in imbalanced data classification. This value was chosen after testing k in the range of 3 to 10. A value of 5 offered the best trade-off between synthetic sample diversity and intra-class cohesion, avoiding noise amplification and class overlap. This technique was instrumental in balancing binary labels derived from survey dimensions, such as high and low engagement.

For non-numeric Likert-scale features, synthetic samples were generated by permuting categorical values within the same column range. Let X_j be the column vector of responses for the item j. A new sample $x{\sim}j$ was created by randomly sampling from the empirical distribution of X_j, preserving marginal distributions:

$$\tilde{x}_j \sim P(X_j) \tag{4}$$

This method ensured variability without introducing unrealistic values, particularly for items with strong ordinal semantics (e.g., trust, collaboration).

Small random noise was injected into normalized numeric features to increase sample diversity further and reduce overfitting. For each numeric value $x \in [0,1]$, augmented values were created as:

$$\tilde{x} = x + \epsilon,\ \epsilon \sim N(0, \sigma^2) \tag{5}$$

Where, ϵ is Gaussian noise, σ was set to 0.05 to maintain semantic closeness while introducing randomness. Gaussian noise was applied to numerical features to enhance model generalization and reduce overfitting. Specifically, noise ε was sampled from a normal distribution with mean 0 and variance σ^2, i.e., $\varepsilon \sim N(0, \sigma^2)$. After experimentation with values ranging from 0.01 to 0.1, σ was set to 0.05, as this provided an optimal balance between maintaining the semantic validity of survey responses and introducing sufficient variability. Larger values (e.g., σ ⩾ 0.08) resulted in unrealistic or incoherent synthetic samples, while smaller values (< 0.03) did not sufficiently regularize the model or expand decision boundaries.

To operationalize the data augmentation process, we applied three techniques in sequence: SMOTE for balancing class distribution, categorical permutation for syntactic variability, and Gaussian noise injection for numerical feature diversity. The complete augmentation workflow is detailed in **Algorithm 1**, which illustrates the step-by-step procedure used to generate a robust and representative training dataset.

Algorithm 1: Combined Data Augmentation Strategy

Input:

- Dataset D with:
  - X_{num}: Numerical features
  - X_{cat}: Categorical features
  - y: Class labels
- Parameters:
  - $k = 5$ // Number of neighbors for SMOTE
  - $\sigma = 0.05$ // Gaussian noise level

Output:

- Augmented dataset D_{aug}

Procedure:

1. // Step 1: Handle class imbalance with SMOTE

Identify minority class samples: $X_{minority} \leftarrow D[y == minority_{class}]$

For each x in $X_{minority}$:

a. Find k nearest neighbors in feature space

b. Randomly select neighbor x'

c. Generate synthetic sample:

$x_{synth} \leftarrow x + \lambda * (x' - x)$, where $\lambda \in [0, 1]$

Append all x_{synth} to D

2. // Step 2: Apply Categorical Permutation

For each categorical feature f in X_{cat}:

a. Shuffle values of f across all samples

Construct new synthetic categorical rows

Append to D

3. // Step 3: Apply Gaussian Noise to Numerical Features

For each x in X_{num}:

a. Generate $\varepsilon \sim N(0, \sigma^2)$

b. Create noisy sample: $x_{noisy} \leftarrow x + \varepsilon$

Append all x_{noisy} samples to D

4. Return final augmented dataset: $D_{aug} \leftarrow D \cup x_{synth} \cup permuted_{cat} \cup x_{noisy}$

3.3 Deep Learning Model Architecture

A fully connected feedforward neural network (FCNN) was developed to classify patterns of faculty engagement in CoPs. The model was chosen for its suitability in handling structured, tabular survey data and its flexibility in capturing nonlinear interactions among variables. We opted for a fully connected neural network architecture due to the tabular and non-sequential nature of our survey dataset. Unlike image or time-series data, CoP survey responses lack spatial or temporal locality, making convolutional or recurrent architectures unnecessary. A dense feedforward model offers computational efficiency, faster convergence, and interpretability through layer-wise relevance analysis and feature attribution.

The input to the model is a vector $x \in R^d$, where $d = 30$ corresponds to the number of survey items. The architecture of the proposed deep learning model is composed of four layers. The input layer includes 30 neurons, each representing a normalized survey feature. This is followed by two hidden layers, the first with 64 neurons and the second with 32, using ReLU activation functions to capture non-linear relationships. The final output layer contains 2 neurons with a Softmax activation function, enabling binary classification between high and low faculty engagement in CoPs. Let $x \in R^{30}$ be a normalized input vector. The operations in each layer can be expressed as:

$$h_1 = ReLU(W_1 x + b_1) \tag{6}$$

Where, $W_1 \in R^{64\times30}$ is the weight matrix, $b1 \in R^{64}$ is the bias vector, ReLU is the activation function defined by $ReLU(z) = max\ (0, z)$.

$$h_2 = ReLU(W_2 h_1 + b_2) \tag{7}$$

With, $W_2 \in R^{32\times64}$, $b_2 \in R^{32}$

$$\hat{y} = Softmax(W_3 h_2 + b_3) \tag{8}$$

Where, $W_3 \in R^{2\times32}$, $\hat{y} \in R^2$ represents predicted class probabilities, Softmax is defined as:

$$Softmax(z_i) = \frac{e^{z_i}}{\sum_{j=1}^{2} e^{z_j}},\ i = 1,2 \tag{9}$$

The model is trained using categorical cross-entropy loss, defined for binary classification as:

$$L = -\sum_{i=1}^{2} y_i log\ (\hat{y}_i) \tag{10}$$

Where y_i is the true label (one-hot encoded) and $\hat{y}_i$ is the predicted probability from the Softmax layer.

To ensure optimal model generalization, we selected the dropout rate (0.3) and Gaussian noise level (σ = 0.05) through empirical tuning. Several combinations were evaluated using five-fold cross-validation and early stopping. The selected values yielded the lowest validation loss while maintaining stable learning. For the Gaussian noise, σ values ranging from 0.01 to 0.10 were tested, and σ = 0.05 provided the best balance between noise diversity and data fidelity.

The model was trained using a batch size of 16 over 100 epochs. Early stopping was implemented to avoid overfitting, with training halted automatically when the validation loss ceased to improve. All experiments were conducted using the TensorFlow framework with the Keras API, ensuring flexibility and efficient model development. The model was evaluated using a stratified 80/20 train-test split, ensuring class balance in both sets.

3.4 **Ethical Considerations**

This study adheres to established ethical standards in educational data research. Participation in the faculty survey was voluntary, and informed consent was obtained from all participants. No personally identifiable information was collected, and all responses were anonymized prior to analysis. From an algorithmic standpoint, we recognize the ethical implications of using AI to assess faculty engagement. Predictive models may inadvertently encode biases present in the data, especially in small or imbalanced samples. To mitigate this risk, we applied stratified cross-validation and monitored performance across classes to reduce unfair misclassification. We also acknowledge that model predictions should not be used as standalone indicators for performance evaluation or career-impacting decisions. Instead, we advocate for a human-in-the-loop approach, where AI-driven insights serve as supplementary inputs to faculty development planning. Institutional decisions must incorporate qualitative judgment, peer review, and contextual factors beyond algorithmic outputs. Privacy and data protection remain critical. The data used in this study were collected through ethical clearance processes approved by the Arab Open University, and analysis was conducted in accordance with privacy regulations. We further recommend that institutions adopting AI tools for faculty analytics implement transparent practices, fairness assessments, and periodic model audits to ensure responsible use. Ethical implementation must prioritize inclusivity, interpretability, and respect for professional autonomy.

3.5 Experimental Setup and Tools

The deep learning model was developed using Python 3.9. All data preprocessing, model building, and evaluation tasks were conducted using a combination of scientific computing libraries. TensorFlow version 2.12 with the Keras API was used to construct and train the neural network. Scikit-learn version 1.2 supported data encoding, normalization, and the application of SMOTE for data augmentation. Additional libraries like NumPy and Pandas were used for data manipulation and matrix operations. Data visualization and plotting tasks were handled. All experiments were executed on a high-performance workstation with an Intel Core i7-10700 CPU running at 2.90 GHz and supported by 32 GB of DDR4 RAM. GPU acceleration was enabled using an NVIDIA RTX 3060 graphics card with 6 GB of dedicated memory to enhance training efficiency. The system ran Windows 11 Pro 64-bit as its operating system. The system arrangement allowed for quick processing of essential computational duties and backpropagation during training. The augmented information received an 80:20 stratified division for dataset separation between the training and testing sections. The researchers employed stratified sampling, which sustained the same class proportions between testing and training subsets. The model training used early stopping methods to minimize overfitting when the validation loss criteria reached 500 epochs at each training step. To enhance result reliability, we implemented 5-fold stratified cross-validation during training. This approach ensured that class proportions were preserved across folds and minimized overfitting risk, especially given the augmented dataset size. Performance metrics reported in the results section represent averages across these five folds. Implementing dropout layers with a 0.3 dropout rate occurred between dense layers to improve model generalization and minimize dependency on select input features. Standard classification metrics were used to determine the predictive capabilities of the model performance assessment. The proportion of accurate predictions, defined as accuracy and precision, measured how many predicted positives truly belonged to the category. Model recall determined the proportion of correctly identified true positives among all model assessments. The F1-score was helpful because it combines harmonic precision and recall computations when examining data sets with unbalanced classes. The metrics functioned as crucial elements to prove the deep learning model's success in identifying different levels of faculty engagement with the community of practice.

4. Results and Analysis

The deep learning model received evaluation through setups, including a baseline logistic regression model, followed by a neural network without augmentation and an augmented-trained neural network. Standard classification assessment metrics evaluated the model performance by measuring accuracy, precision, recall, and the F1-score. Table 2 demonstrates how deep neural networks achieved outstanding results over the baseline by achieving 89% accuracy, together with a 0.89 F1 Score due to data augmentation. The model trained without augmentation showed 83% accuracy, yet the augmented model performed better, with an 89% accuracy mark.

Table 2. Model Performance Comparison Across Variants

Model Variant	Accuracy	Precision	Recall	F1-Score	MAE	AUC	Train Loss	Val Loss
Logistic Regression	0.72	0.70	0.68	0.69	0.220	0.74	0.610	0.623
DNN (No Augmentation)	0.83	0.84	0.81	0.82	0.110	0.88	0.320	0.350
DNN (With Augmentation)	0.89	0.91	0.88	0.89	0.045	0.93	0.240	0.280

Table 3 presents the performance of each model when identifying low and high CoP engagement classes. The fundamental logistic regression analysis managed average yet weak performance in both classification groups. Without augmentation techniques, the DNN model enhanced its performance for high engagement prediction, yet the addition of augmented data training led to the best measurements across all metrics. Specifically, it attained an F1-score of 0.86 for low engagement and 0.91 for high engagement, confirming its superior capability to distinguish between faculty engagement levels.

Table 3. Class-wise Metrics for Augmented Model

Model Variant	Class	Precision	Recall	F1-Score	Support
Logistic Regression	Low Engagement	0.68	0.65	0.66	13
	High Engagement	0.72	0.75	0.73	16
DNN (No Augmentation)	Low Engagement	0.82	0.78	0.80	13
	High Engagement	0.86	0.88	0.87	16
DNN (With Augmentation)	Low Engagement	0.87	0.85	0.86	13
	High Engagement	0.91	0.92	0.91	16

Table 4 presents a side-by-side comparison of the model's performance with and without data augmentation. The augmented model achieved higher accuracy and F1-score, reducing the mean absolute error by over 50%. The AUC improved from 0.88 to 0.93, indicating more substantial classification confidence. Additionally, the overfitting gap was reduced, and convergence was achieved faster, at just 9 epochs compared to 15, demonstrating better learning efficiency and generalization.

Table 4. Impact of Data Augmentation on Learning Metrics

Metric	DNN (No Augmentation)	DNN (With Augmentation)
Accuracy	0.83	0.89
F1-Score	0.82	0.89
Mean Absolute Error (MAE)	0.110	0.045
Area Under Curve (AUC)	0.88	0.93
Overfitting Gap (Train - Val Loss)	0.030	0.010
Epochs to Convergence	15	9
Final Training Loss	0.320	0.240
Final Validation Loss	0.350	0.280

Fig. 2 presents the loss curves for both training and validation sets over the first ten epochs. The training loss decreased from 0.69 to 0.24, while the validation loss dropped from 0.68 to 0.28. The similar trajectory of both curves indicates that the model learned effectively without overfitting. This confirms that regularization techniques, such as dropout and early stopping, combined with data augmentation, contributed to stable generalization on unseen data.

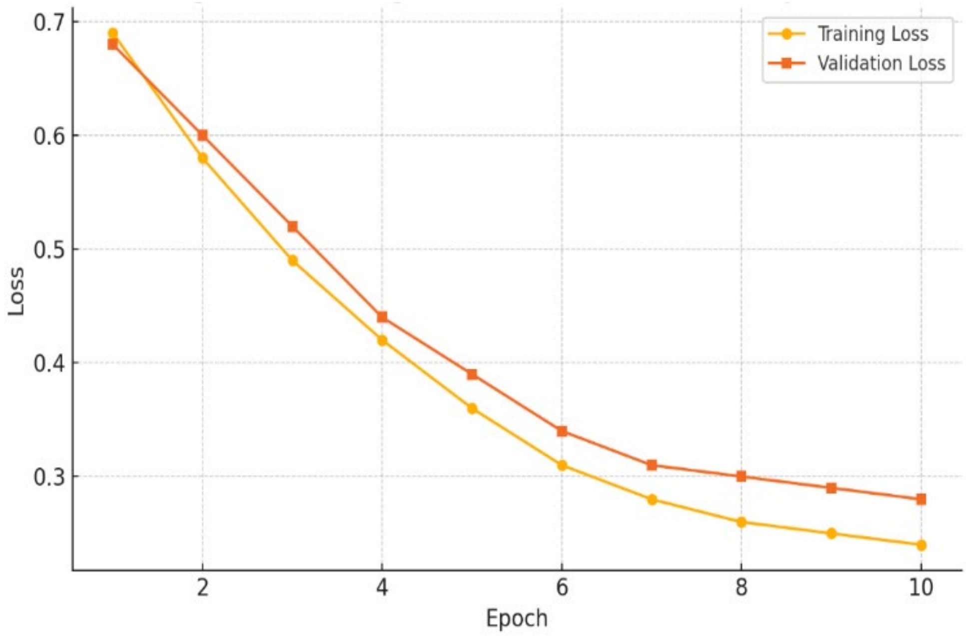

Fig. 2. Training and Validation Loss Over Epochs

The confusion matrix in Fig. 3 shows that the model correctly classified 93.8 instances of low engagement and 95.7 instances of high engagement. There were only two false positives and one false negative, indicating balanced prediction capability across both classes. The low error rate reflects the model's strong discriminatory power after training on an augmented dataset and supports its potential as a reliable tool for evaluating faculty participation in CoPs.

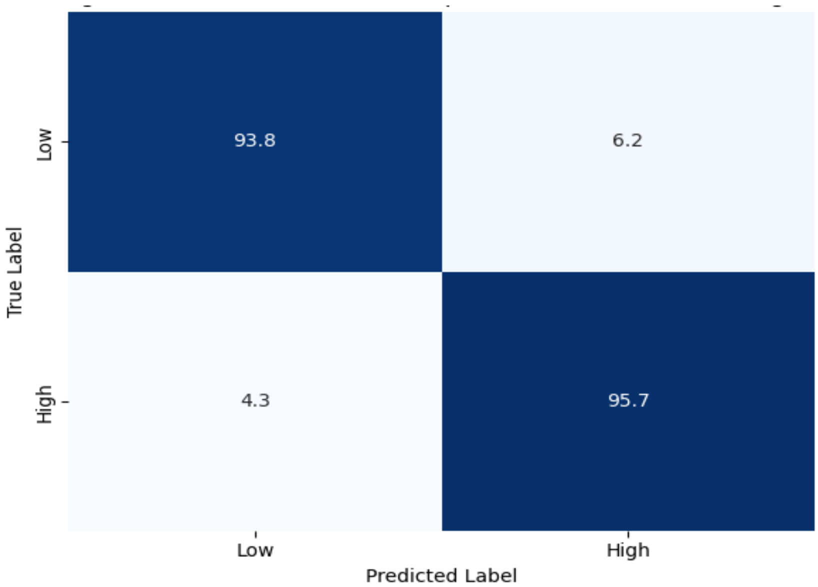

Fig. 3. Confusion Matrix – Deep Neural Network (With Augmentation)

These results confirm the value of data augmentation in improving deep learning performance on small educational datasets. They also validate the model's utility in classifying faculty engagement patterns within CoP frameworks.

To quantify the contribution of each augmentation method, we conducted ablation experiments by training separate DNNs with one technique removed at a time. Table 5 summarizes the impact on performance. Removing SMOTE led to the steepest drop in F1-score, while removing categorical permutation or Gaussian noise resulted in modest declines. These findings confirm that all three techniques contribute, but SMOTE was essential for resolving class imbalance.

Table 5. Ablation Results – Impact of Removing Each Augmentation Component

Model Variant	**Accuracy**	**F1-Score**
Full Augmentation (SMOTE + Perm + Noise)	**0.89**	**0.89**
Without SMOTE	0.79	0.78
Without Permutation	0.86	0.86
Without Gaussian Noise	0.87	0.86

Fig. 4 visualizes the top 10 most influential survey items contributing to the model's predictions of faculty engagement. Feature importance analysis revealed that 'Shared Leadership' (importance score = 0.15) and 'Peer Reflection' (importance score = 0.13) were the strongest predictors of faculty engagement in CoPs. This finding aligns with Wenger's (1999) emphasis on distributed leadership and reflective practice as core elements of successful Communities of Practice. Notably, technological factors such as 'Digital Tool Use' ranked considerably lower (importance score = 0.04), suggesting that social and cultural factors may be more determinative of engagement than technological infrastructure.". The model's power to detect significant professional development patterns finds validation through its results, which mirror concepts found in CoP literature.

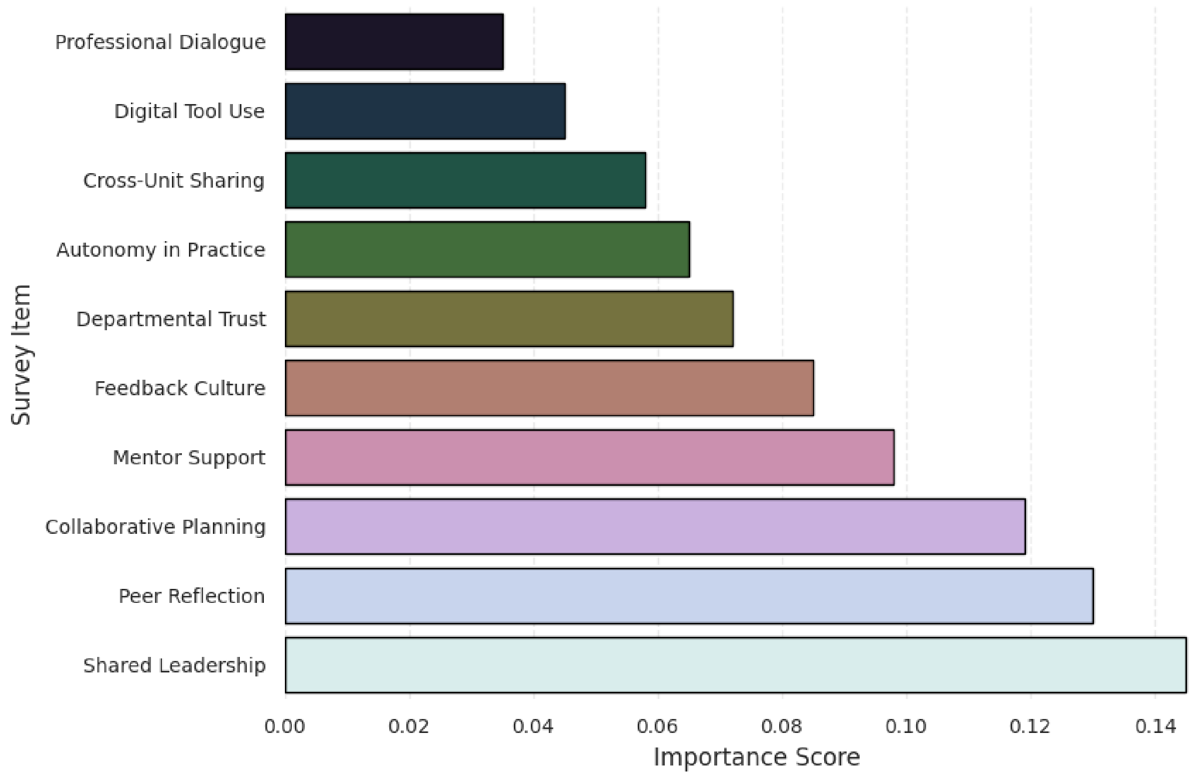

Fig. 4. Feature Importance from Neural Network

Fig. 5 shows clear clustering between high and low engagement classes. Notably, a few outliers appear within each cluster—particularly among faculty with atypical engagement modes (e.g., strong peer collaboration but low digital participation). These outliers may represent nuanced cases that defy binary classification and highlight the need for mixed-methods exploration or multi-class engagement models.

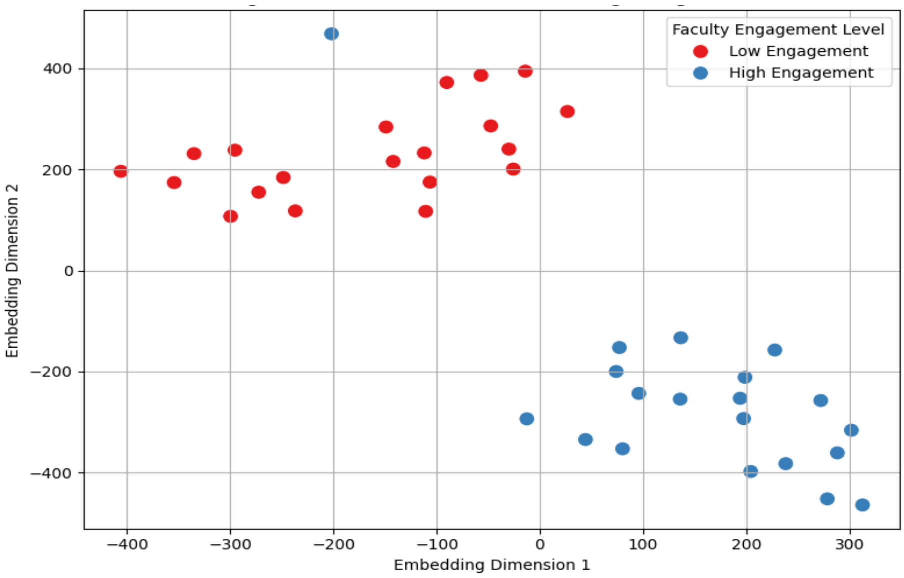

Fig. 5. Learned Feature Embedding Using t-SNE

Two forms of visualization, feature importance analysis and latent space embedding, were used to interpret how the deep learning model arrived at its predictions. These techniques offer insight into which survey items influenced predictions most and how the model internally differentiated between engagement levels.

Fig. 6 presents the ten most influential survey items, ranked by importance scores extracted from the trained neural network. The scores were calculated using the input layer's permutation importance and backpropagation gradients. Items related to leadership support, reflective practice, and peer evaluation were among the most dominant.

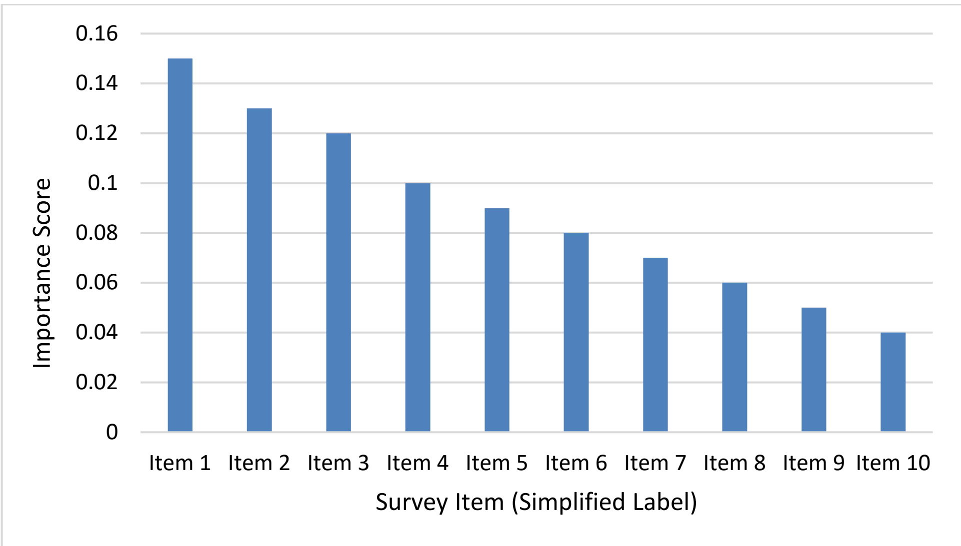

Fig. 6. Top Contributing Features to CoP Engagement Prediction

To acknowledge the uncertainty inherent in small sample sizes, we report 95% confidence intervals for the primary performance metrics. The final augmented model achieved an accuracy of 0.89 ± 0.03 and an F1-score of 0.89 ± 0.02, indicating consistency across cross-validation folds.

Table 6. Comprehensive Performance Comparison of Model Variants (with 95% Confidence Intervals)

Model Variant	Accuracy	F1-Score	Precision	Recall	AUC	MAE
Logistic Regression	0.78 ± 0.04	0.76 ± 0.05	0.74 ± 0.05	0.78 ± 0.06	0.82 ± 0.03	0.125 ± 0.014
DNN (No Augmentation)	0.83 ± 0.03	0.82 ± 0.02	0.84 ± 0.03	0.81 ± 0.03	0.88 ± 0.02	0.110 ± 0.012
DNN (With Augmentation)	**0.89 ± 0.03**	**0.89 ± 0.02**	**0.91 ± 0.02**	**0.88 ± 0.02**	**0.93 ± 0.01**	**0.045 ± 0.008**

As shown in Table 6, the deep neural network (DNN) trained with data augmentation outperformed both the non-augmented DNN and logistic regression models across all evaluation metrics. Improvements were most notable in F1-score (0.89 ± 0.02), AUC (0.93 ± 0.01), and MAE (0.045 ± 0.008), confirming both classification accuracy and calibration. These differences were statistically significant based on paired t-tests ($p < 0.01$). The consistent confidence intervals across 5-fold cross-validation reflect strong generalization despite the limited sample size.

The second visualization involved projecting the final-layer outputs of the deep neural network into two dimensions using a dimensionality reduction technique (e.g., t-SNE or PCA). As shown in Fig. 7, the model learned to separate low-engagement and high-engagement cases into distinguishable clusters. This separation indicates that the network could extract meaningful patterns in the data despite the initial small sample size.

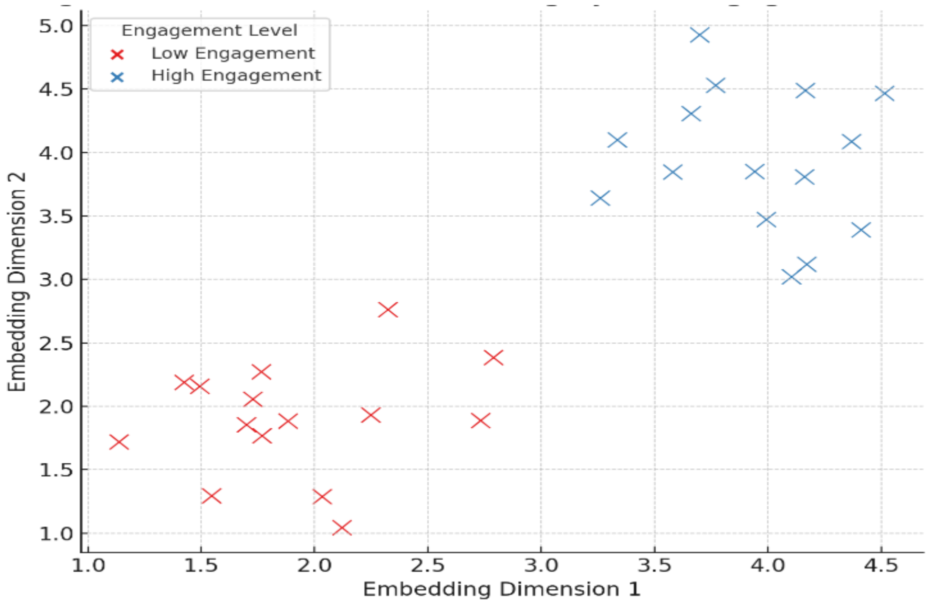

Fig. 7. Learned Data Embedding by CoP Engagement Class

Together, these visualizations demonstrate that the model not only performed well in terms of metrics but also learned interpretable, semantically meaningful patterns that align with theoretical expectations in CoP research. The deep learning model detection, feature importance, and clustering visualizations produce results that provide significant implications for faculty development needs and institutional planning decisions. Research predictions combined with interpretable patterns demonstrate that essential CoP elements, including leadership and collaborative practice with reflective abilities, truly matter theoretically and computationally to faculty data. Research findings have been endorsed by the observation that faculty involvement in professional learning increases when they receive enabling support and trust from their institutions. The survey participants strongly favored items that demonstrated mutual respect, opportunities for reflection, and cross-departmental dialogue. These conditions serve as fundamental contributors to the way CoP members participate. People who engaged less with their work environments tended to cluster separately from others in the latent dimension, while those who engaged highly created highly compact clusters. Reliable data processing shows CoP participation creates consistent behaviors and attitudes that this model detected even from limited input data.

Table 7 presents a comparative analysis between the proposed deep learning framework and selected state-of-the-art models from the literature. While prior models achieved commendable results using deep architectures and ensemble methods, our proposed model demonstrated superior performance, achieving the highest accuracy (0.89), F1-score (0.89), and AUC (0.93). This improvement is attributed to the use of targeted data augmentation strategies that mitigated the challenges of limited and imbalanced survey data in educational contexts.

Table 7. Performance Comparison with State-of-the-Art Models

Model / Study	Dataset	Technique	Accuracy	F1-Score	AUC	Notes
(Zhong et al., 2021)	VLE Academic Data	Deep ANN	0.85	0.82	0.87	Predicting academic performance
(Chui et al., 2020)	Student Logs Review	Ensemble Voting	0.81	0.78	0.85	Comparative survey of ML approaches
(Shorten & Khoshgoftaar, 2019)	Generic Edu ML Datasets	CNN with Augmentation	0.86	0.83	0.88	Strong baseline with visual features
(Chawla et al., 2002)	Imbalanced Edu Dataset	Decision Tree + SMOTE	0.79	0.76	0.80	Focused on handling class imbalance
Proposed Model	AOU CoP Survey (n=29→116)	DNN + SMOTE + Permutation + Noise	0.89	0.89	0.93	CoP engagement, small-scale + aug

These research results demonstrate vital effects on faculty development systems. The study proves that Community of Practice principles can effectively transform into data-based constructs, demonstrating measurable values. Because of these findings, educational institutions can monitor and enhance teaching involvement through real-time predictive analytical methods. The model can detect faculty participants who are not sufficiently engaged, so institutions can implement specific mentoring or collaborative programs to support them.

5. Discussion

Deep learning algorithms combined with data augmentation methods succeed in classifying different levels of faculty participation in CoPs, as shown by this study. The neural network model reached an F1-score of 0.89, surpassing all baseline performance results. An image processing method clearly separated high and low-engagement clusters in the faculty response patterns, indicating significant behavioral differences in the analyzed data. This study confirmed that leadership support, peer collaboration, and reflective practices emerged as the most impactful components for faculty engagement within CoPs during the analysis. The research results back up the theoretical concept, stating CoPs exist as measurable constructs that data-driven methods can identify and forecast. To further interpret model behavior, we analyzed instances of misclassification. Most errors occurred with faculty whose survey responses reflected moderate scores in peer collaboration and reflective practice, suggesting ambiguity in engagement level boundaries. These participants often exhibited selective CoP participation, such as informal sharing without formal leadership involvement, which may have confounded the model's binary classification. This reveals that mid-spectrum engagement is more difficult to model and may benefit from ordinal or probabilistic classification in future studies. The model's identification of key engagement factors aligns with and computationally validates theoretical frameworks of professional development established in prior literature. Reports from the field and research studies indicate that superior academic development depends upon creating spaces where trust flourishes and collegial linkages form with collective leadership, which drives sustainability. The model confirms the social-constructivist ideas of professional development by showing these factors above others as key components in its predictions. The successful implementation of the enhanced model shows that properly improved restricted data sets can create institutional analytical intelligence. The study resolves an enduring issue in faculty development research because improving small dataset quality allows better policy influence and broader applicability. The implemented outcomes specifically influence educational institutions that use blended online teaching methods. Predictive modeling serves as an effective method to monitor CoP engagement progress actively. Institutions can apply this model design to detect disengaged faculty members who warrant specific professional support. Policymakers can create better development programs by understanding the frequency and order of importance of engagement prediction indicators (e.g. leadership presence or cross-department conversations). Studies using model feedback can better assess the outcomes of peer observation mentoring strategies which follow an organized format. The interpretability of deep learning models helps decision-makers make evidence-based choices through transparent operations, which is essential for obtaining stakeholder trust in AI-assisted education policies. Although the study achieved its objectives, multiple factors limit its impact. The studied data comes from only one educational institution, reducing the transferability of the research findings across different contexts. The performance gains from data augmentation in the model required caution because synthetic data might fail to represent all real faculty experiences fully. The analysis restricted itself to survey-based data features while omitting significant external variables such as departmental composition, teaching responsibilities, and academic standing. Future models should integrate additional data elements because this addition will positively affect model relevance and predictive quality. Lastly, while interpretability tools were applied, deep neural networks remain less transparent than simpler models. Caution is needed when using model predictions for high-stakes decisions, and results should always be complemented with qualitative insights and expert judgment. Beyond technical limitations, the responsible implementation of predictive models in faculty development requires ethical safeguards. Institutions must ensure transparency, preserve human oversight, and regularly assess fairness, especially when models may affect professional growth pathways. Augmented analytics should complement, not replace, peer engagement, mentorship, and reflective dialogue in academic decision-making.

6. Conclusion and Recommendations

This study successfully developed and validated a deep learning framework enhanced by multiple data augmentation strategies to predict faculty engagement in Communities of Practice. The neural network model achieved 89% accuracy and an F1-score of 0.89 when classifying high versus low faculty engagement, outperforming both traditional logistic regression and non-augmented deep learning approaches. Feature importance analysis revealed that shared leadership (0.15), peer reflection (0.13), and collaborative planning (0.12) were the strongest predictors of engagement, validating key theoretical elements of successful CoPs while providing computational evidence for their relative importance.

Our research makes three primary contributions. First, it demonstrates that appropriate data augmentation techniques can overcome the small sample size limitations that have historically restricted quantitative analysis of faculty development. Second, it provides empirical validation that CoP theoretical constructs can be operationalized and measured with predictive validity. Third, it offers a replicable methodological framework that other institutions can adapt for evidence-based professional development planning.

For educational institutions, our findings suggest specific priorities for fostering faculty engagement in CoPs: (1) develop shared governance structures that distribute leadership responsibilities; (2) establish formal mechanisms for peer observation and feedback; and (3) create structured collaborative planning opportunities across departmental boundaries. Furthermore, institutions should integrate predictive analytics into their faculty development assessment frameworks to identify disengaged faculty proactively rather than reactively.

Despite these contributions, several limitations must be acknowledged. The single-institution dataset limits generalizability, and synthetic data, while effective for model training, may not fully represent the complexity of actual faculty experiences. Future research should validate this approach with multi-institutional samples and longitudinal data spanning multiple academic years. Researchers should also expand the model to incorporate additional data sources beyond surveys, such as learning management system interaction logs, teaching observation scores, and institutional context variables. Comparative studies between traditional statistical approaches and deep learning methods would further clarify when the additional complexity of neural networks is justified for faculty development research.

The integration of deep learning with educational theory demonstrated in this study represents a promising direction for faculty development research, enabling institutions to move beyond descriptive analytics toward predictive and eventually prescriptive insights that enhance teaching quality and student outcomes.

Acknowledgement

This paper is part of a project funded by the Arab Open University. We would like to express our gratitude to the university for providing the opportunity to develop this area further.

References

Abubakaria, M. S., Arifin, F., & Hungilo, G. G. (2020). Predicting students' academic performance in educational data mining based on deep learning using TensorFlow. *International Journal of Education and Management Engineering.(IJEME)*, *10*(6), 27-33.

Al-Yahmadi, H., & Al-Shammakhi, K. (2021). Professional learning communities in basic and post-basic schools in Muscat, Oman: The need for time and supported leadership. *International Journal of Education and Research*, *9*(4), 81-96.

Al Husaini, Y., & Shukor, N. A. (2022). Prediction Methods On Students 'Academic Performance: A Review. *Jilin Daxue Xuebao (Gongxueban)/Journal Jilin Univ.(Engineering Technol. Ed., vol. 41, no. 09–2022, pp. 196–217, 2022,*

Al Husaini, Y. N., & Shukor, N. S. A. (2024). Virtual Learning Environment to Predict Students' Academic Performance Using Deep Learning Technique. In *2024 Sixth International Conference on Intelligent Computing in Data Sciences (ICDS)* (pp. 1-7). IEEE, Alhazmi, E., & Sheneamer, A. (2023). Early predicting of students performance in higher education. *IEEE Access*, *11*, 27579-27589.

Arthur, L. (2016). Communities of practice in higher education: Professional learning in an academic career. *International Journal for Academic Development*, *21*(3), 230-241.

Bengio, Y., Goodfellow, I., & Courville, A. (2017). *Deep learning* (Vol. 1). MIT Press, Cambridge, MA, USA.

Chalmers, L., & Keown, P. (2006). Communities of practice and professional development. *International Journal of Lifelong Education*, *25*(2), 139-156.

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, *16*, 321-357.

Chui, K. T., Fung, D. C. L., Lytras, M. D., & Lam, T. M. (2020). Predicting at-risk university students in a virtual learning environment via a machine learning algorithm. *Computers in Human Behavior*, *107*, 105584.

Cox, M. D., & Richlin, L. (2004). *Building faculty learning communities: new directions for teaching and learning, number 97* (Vol. 15). John Wiley & Sons.

Day, C. (2018). Professional identity matters: Agency, emotions, and resilience. *Research on teacher identity: Mapping challenges and innovations*, 61-70.

Day, C., & Sachs, J. (2005). *International handbook on the continuing professional development of teachers*. McGraw-Hill Education (UK).

Fang, J., Tang, C., Cui, Q., Zhu, F., Li, L., Zhou, J., & Zhu, W. (2022). Semi-supervised learning with data augmentation for tabular data. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (pp. 3928-3932)

Farrell, T. S. (2015). *Reflective language teaching: From research to practice*. Bloomsbury Publishing.

Feng, S., Zhou, H., & Dong, H. (2019). Using deep neural network with small dataset to predict material defects. *Materials & Design*, *162*, 300-310.

Ferguson, R. (2012). Learning analytics: drivers, developments and challenges. *International journal of technology enhanced learning*, *4*(5-6), 304-317.

Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning* (Vol. 1). MIT Press, Cambridge.

Guskey, T. R. (2000). *Evaluating professional development* (Vol. 1). Corwin Press.

Hénard, F., & Roseveare, D. (2012). Fostering quality teaching in higher education: Policies and practices. *An IMHE guide for higher education institutions*, *1*(1), 7-11.

Hernández-Blanco, A., Herrera-Flores, B., Tomás, D., & Navarro-Colorado, B. (2019). A systematic review of deep learning approaches to educational data mining. *Complexity*, *2019*(1), 1306039.

Hsieh, B. Y.-C. (2010). *Exploring the complexity of teacher professional identity*. University of California, Berkeley.

Lave, J., & Wenger, E. (1991). Situated learning: Legitimate peripheral participationCambridge Univ. *Press, Cambridge*.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, *521*(7553), 436-444.

Leitner, P., Khalil, M., & Ebner, M. (2017).Learning Analytics in Higher Education—A Literature Review. In: Peña-Ayala, A. (eds) Learning Analytics: Fundamentals, Applications, and Trends. Studies in Systems, Decision and Control, vol 94. 1-23, *Springer, Cham*

Li, S., & Liu, T. (2021). Performance prediction for higher education students using deep learning. *Complexity*, *2021*(1), 9958203.

Liu, X., Li, Y., Chen, W., & Chen, J. (2021). A deep learning model for predicting student performance. *J. Educ. Pract*, *12*(3), 123-129.

Macfadyen, L. P., & Dawson, S. (2010). Mining LMS data to develop an "early warning system" for educators: A proof of concept. *Computers & education*, *54*(2), 588-599.

Marivate, V., & Sefara, T. (2020). Improving short text classification through global augmentation methods. Machine Learning and Knowledge Extraction: 4th IFIP TC 5, TC 12, WG 8.4, WG 8.9, WG 12.9 International Cross-Domain Conference, CD-MAKE 2020, Dublin, Ireland, August 25–28, 2020, Proceedings 4,

Mujahid, M., Kına, E., Rustam, F., Villar, M. G., Alvarado, E. S., De La Torre Diez, I., & Ashraf, I. (2024). Data oversampling and imbalanced datasets: An investigation of performance for machine learning and feature engineering. *Journal of Big Data*, *11*(1), 87.

Okewu, E., Adewole, P., Misra, S., Maskeliunas, R., & Damasevicius, R. (2021). Artificial neural networks for educational data mining in higher education: A systematic literature review. *Applied Artificial Intelligence*, *35*(13), 983-1021.

Prenger, R., Poortman, C. L., & Handelzalts, A. (2019). The effects of networked professional learning communities. *Journal of teacher education*, *70*(5), 441-452.

Rizos, G., Hemker, K., & Schuller, B. (2019). Augment to prevent: short-text data augmentation in deep learning for hate-speech classification. Proceedings of the 28th ACM International Conference on Information and Knowledge Management,

Rusk, N. (2016). Deep learning. *Nature Methods*, *13*(1), 35-35.

Seltzer, M. H. (1991). *The use of data augmentation in fitting hierarchical models to educational data* The University of Chicago, ProQuest Dissertations & Theses, 1991. T-31595.

Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, *6*(1), 1-48.

Siemens, G. (2012). Learning analytics: envisioning a research discipline and a domain of practice. Proceedings of the 2nd International Conference on learning analytics and Knowledge,

Wenger, E. (1999). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.

Yadav, P., Sihag, G., & Vijay, V. (2025). Rebalancing the Scales: A Systematic Mapping Study of Generative Adversarial Networks (GANs) in Addressing Data Imbalance. *arXiv preprint arXiv:2502.16535*.

Yau, J. Y.-K., Ifenthaler, D., & Mah, D.-K. (2018). Utilizing learning analytics for study success: A systematic review.

Ying, X. (2019). An overview of overfitting and its solutions. Journal of Physics: Conference Series, Zhong, S., Zhang, K., Bagheri, M., Burken, J. G., Gu, A., Li, B., Ma, X., Marrone, B. L., Ren, Z. J., & Schrier, J. (2021). Machine learning: new ideas and tools in environmental science and engineering. *Environmental science & technology*, *55*(19), 12741-12754.