

## Explainable Multi-View Modeling of AI-Driven Personalized Learning Adoption in TVET Systems

Xiao Zhang<sup>1,2\*</sup>, Jinghong Low<sup>1</sup>

<sup>1</sup>Faculty of Management, Multimedia University, Selangor, Malaysia, 63100

<sup>2</sup>School of Law, History and Culture, Yulin University, Yulin, Shaanxi, China, 719000

*zhang8xiao89@163.com*

**Abstract.** Artificial intelligence-driven personalized learning is increasingly embedded within Technical and Vocational Education and Training (TVET) systems as a data-enabled service innovation. However, large-scale institutional adoption remains uneven due to socio-technical, governance, and organizational constraints. This study conceptualizes TVET personalization as a digital service system and investigates the critical determinants of adoption through an integrated explanatory-predictive framework. Building on UTAUT2 and organizational readiness perspectives, we model personalization quality, trust in AI, transparency, perceived privacy risk, instructor support, and facilitating conditions as key service-system drivers of behavioral intention and sustained use. Recent empirical studies on Generative AI adoption in higher education further validate the applicability of the UTAUT2 framework, confirming that facilitating conditions and performance expectancy are critical determinants for integrating educational technologies (Papadakis et al., 2025). We propose a Theory-Guided Multi-View Multi-Task (TG-MVMT) framework that fuses survey-based latent constructs with learning-platform behavioral traces and incorporates institution-aware regularization to enhance cross-campus robustness. The model is benchmarked against strong predictive baselines (logistic regression, random forest, gradient boosting, and TabNet) under nested cross-validation, and interpretability is achieved using SHAP-based attribution to translate predictive signals into actionable governance levers. Results indicate that combining perceptual factors with early behavioral indicators significantly improves adoption prediction and reveals a concentrated set of dominant determinants, with personalization quality and trust in AI consistently emerging as primary drivers. The findings contribute to service informatics research by demonstrating how explainable multi-view modeling can support AI governance, instructor-in-the-loop orchestration, and scalable deployment across heterogeneous TVET institutions.

**Keywords:** TVET service systems; AI-driven personalization; service informatics; technology adoption; explainable machine learning; digital transformation

## 1. Introduction

AI-driven personalized learning in TVET is not only a technical integration problem but also a socio-technical adoption challenge. Prior research on educational technology uptake has largely relied on individual-level acceptance models such as the technology acceptance model (Davis et al., 1989) and the unified theory of acceptance and use of technology (Venkatesh et al., 2012; DeLone et al., 2003), complemented by system-level success perspectives that emphasize information quality, service quality, and net benefits. Empirical comparisons using the SERVQUAL framework further confirm that generative AI significantly outperforms traditional models in responsiveness, assurance, and empathy dimensions, although human oversight remains essential to mitigate hallucination risks (Eirena & Shah, 2025). For institutionally embedded innovations in TVET, organizational and environmental readiness considerations are equally important, as emphasized in the technology-organization-environment framework (VanLehn et al., 2011). Recent studies utilizing the TOE framework also highlight that technological factor, such as relative advantage, complexity, and compatibility, are critical antecedents determining the successful adoption of generative AI technologies (Twaissi et al., 2025). However, these frameworks were mostly validated in contexts where technologies are relatively transparent and their risks are well understood; A systematic review of conversational AI in education also highlights that despite growing adoption, significant research gaps remain regarding cultural adaptation, infrastructural disparities, and ethical data governance frameworks (Alkishri et al., 2025). AI personalization introduces new governance concerns (data protection, explainability, accountability) that can reshape adoption dynamics. Yet TVET differs from general academic education because learning is competency-based, time-constrained, and tightly coupled with workshop practice and equipment availability. Consequently, the adoption decision depends not only on perceived learning benefits but also on whether the AI recommendations fit practical skill pathways, assessment regimes, and instructor workflow. Learning analytics and educational data mining provide the methodological foundation for capturing these realities through trace data and for evaluating the impact of adaptive interventions (Siemens et al., 2013). Learning analytics has also emerged as a discipline that connects data, models, and organizational capability building, highlighting that analytics value materializes only when institutions develop the processes and expertise to act on insights (Viberg et al., 2022). These insights motivate TVET-specific adoption research that links individual acceptance with institutional analytics readiness. Recent reviews on multimodal learning analytics (MMLA) also highlight that while fusing data sources offers rich perspectives, a critical challenge remains in deeply integrating these data with established learning theories to avoid the interpretability limitations of purely data-driven approaches (Giannakos & Cukurova, 2023).

Privacy and explainability are particularly salient in TVET, where learners often transition directly into employment and may be sensitive to how performance data are used. Privacy has been widely reported as a key obstacle to scaling learning analytics and can undermine willingness to use AI-enabled learning services when governance arrangements are unclear (Khosravi et al., 2022). Recent implementations demonstrate that hybrid architectures combining local large language models with semantic retrieval can serve as a critical security control to mitigate data sovereignty risks and ensure sustained compliance (Eirena & Shah, 2025). At the same time, explainable AI is increasingly viewed as a prerequisite for trustworthy educational AI because stakeholders need to understand why specific learning-path recommendations or risk flags are produced and how they relate to evidence (Khosravi et al., 2022). Therefore, adoption factors in TVET should explicitly incorporate AI trust, transparency, and perceived privacy risk, rather than treating them as secondary concerns.

Methodologically, many adoption studies in education rely solely on cross-sectional surveys and report intention rather than verified use. This approach is insufficient for AI personalization because actual uptake is shaped by interaction patterns (e.g., recommendation acceptance, persistence, time-on-task) that are observable in learning-platform logs. Public learning analytics datasets demonstrate the feasibility of combining behavioral traces with contextual variables at scale (Kuzilek et al., 2017), but

TVET adoption research has rarely leveraged such multi-source evidence. In addition, theory-testing with structural equation modeling requires rigorous assessment of measurement quality and discriminant validity to avoid inflated or misleading path estimates (Fornell et al., 1981). Recent advances provide more reliable discriminant validity criteria for variance-based SEM, which is common in prediction-oriented adoption research (Henseler et al., 2015). These considerations motivate an integrated design that (i) validates constructs with SEM and (ii) benchmarks predictive performance with modern machine-learning models.

To close the above gaps, this study makes three contributions. First, we propose a TVET-specific adoption model that unifies classical acceptance constructs with AI governance and TVET operational conditions. Second, we introduce TG-MVMT, an explainable theory-guided multi-view multi-task predictor that fuses survey-based latent factors with behavioral log features and is evaluated against strong baselines including random forests (Breiman et al., 2001), gradient-boosted trees (Chen et al., 2016), and a competitive deep tabular model (Arik et al., 2021). Third, we operationalize explainability with post-hoc attribution methods to translate predictive signals into decision-relevant determinants for leaders and instructors (Lundberg et al., 2017, Ribeiro et al., 2016). This dual explanatory-predictive strategy complements recent work on learning analytics tool adoption (Gašević et al., 2015) and evidence from vocational learners' adoption of large-scale online learning services (Ifenthaler et al., 2015), and is designed to support scaling AI personalization across heterogeneous TVET providers.

TVET adoption contexts also pose distinct implementation frictions that are often underrepresented in general higher-education adoption studies. Training schedules are frequently block-based, learning occurs across mixed modalities (online, classroom, workshop, and workplace), and instructors must balance formative feedback with safety supervision. As a result, facilitating conditions and instructor support are not merely moderators but can become binding constraints that determine whether AI personalization is perceived as feasible in practice. In addition, TVET learners may include apprentices and adult upskilling participants whose opportunity costs and time pressure are higher than those of traditional students. These realities motivate the explicit modeling of pedagogical compatibility (fit to competency-based pathways) and the evaluation of transfer across institutions with different resource profiles and platform maturity.

Accordingly, the study is driven by three research questions: RQ1: Which psychological, institutional, and AI-governance determinants most strongly explain behavioral intention and sustained use of AI-driven personalization in TVET? RQ2: Can a theory-guided, explainable multi-view predictor achieve higher out-of-sample accuracy than standard SEM and strong tabular baselines while remaining interpretable? RQ3: Do the identified determinants and predictive models generalize across institutions with heterogeneous readiness, and what design choices improve robustness under domain shift?

## **2. Proposed Model: Theory-Guided Multi-View Multi-Task Learning (Tg-Mvmt)**

### **2.1. Study Area**

The empirical setting targeted TVET institutions that deliver competency-based programs combining classroom instruction, workshop practice, and workplace-aligned assessment. The study area is characterized by heterogeneous institutional capacity: some campuses have mature learning management systems, stable broadband, and dedicated instructional designers, whereas others operate under equipment constraints, high instructor workloads, and limited analytics support. This heterogeneity is typical of national or regional TVET systems where campuses serve different local industries (e.g., manufacturing, industrial automation, logistics, hospitality, and health services) and where learners span diverse age groups and prior work experience.

AI-driven personalized learning was piloted within skill modules that include both conceptual knowledge and procedural practice (e.g., safety protocols, machine setup, quality inspection, or process

troubleshooting). The personalized learning component was positioned as an augmentation to existing instruction rather than a replacement, and instructors retained authority to modify learning paths to match workshop schedules and equipment availability. To ensure responsible deployment, the pilot design assumed institution-level governance mechanisms for consent, access control, and data retention, and emphasized transparent communication of how learner data would be used for personalization and analytics.

Because TVET providers frequently aim to scale innovations beyond a single campus, the evaluation design explicitly considers cross-institution transfer and domain shifts. Differences in curriculum sequencing, assessment artifacts, and instructor support structures were treated as realistic sources of distribution shift, motivating robustness checks such as leave-one-institution-out testing and institution-aware modeling.

## 2.2. Ai-Driven Personalized Learning Intervention

The AI-driven personalized learning system provides individualized learning-path recommendations, adaptive practice, and formative feedback. The personalization engine maintains a learner knowledge state and selects the next learning object to maximize expected mastery gains while controlling cognitive load. To support adoption in high-stakes training environments, the system includes an explanation panel (why-this-item) and transparent controls that allow learners and instructors to view, override, and provide feedback on recommendations. Privacy-by-design defaults are enforced (data minimization, role-based access, consent-driven analytics, and configurable retention).

For reproducibility, the personalization policy can be implemented as a contextual bandit on top of a knowledge tracing model. Let  $s_t$  denote the learner state after interaction  $t$ ,  $a_t$  the selected learning object,  $o_t$  the learner response (correctness, latency, hint usage, self-reported confidence), and  $r_t$  the observed learning gain (e.g., delta mastery or quiz improvement). The policy  $\pi(a_t|s_t)$  chooses  $a_t$  to maximize expected cumulative reward subject to exploration.

$$\max_{\pi} \mathbb{E} \left[ \sum_{t=1}^T r_t \right] \quad \text{s.t. } a_t \sim \pi(\cdot | s_t) \\ s_{t+1} = f(s_t, a_t, o_t)$$

## 2.3. Tg-Mvmt

To identify key adoption factors while achieving strong predictive validity, we propose TG-MVMT, a theory-guided multi-view multi-task model. The model fuses (i) survey-based latent factors and (ii) behavioral log features, and simultaneously predicts adoption intention and sustained adoption behavior. Unlike purely correlational models, TG-MVMT injects directional constraints from established adoption theory to improve robustness and interpretability, and includes institution embeddings to capture hierarchical heterogeneity across TVET providers.

Let  $x_i^S$  in  $R^{d_S}$  denote the survey latent-factor vector for participant  $i$  and  $x_i^L$  in  $R^{d_L}$  denote the log feature vector. Two encoders map each view to a shared embedding space:  $h_i^S = f_S(x_i^S)$  and  $h_i^L = f_L(x_i^L)$ , where  $f_S$  and  $f_L$  are multilayer perceptrons with layer normalization and dropout. A gated fusion produces the joint representation.

$$g_i = \sigma(W_g[h_i^S; h_i^L] + b_g), \quad h_i = g_i \odot h_i^S + (1 - g_i) \odot h_i^L$$

An institution embedding  $e_{u(i)}$  is learned for each institution  $u(i)$  of participant  $i$ , and concatenated with  $h_i$  to form the final representation  $r_i = [h_i; e_{u(i)}]$ . Two task-specific heads output probabilities for intention and behavior.

$$\hat{y}_i^{(I)} = \sigma(w_I^T r_i + b_I), \quad \hat{y}_i^{(B)} = \sigma(w_B^T r_i + b_B).$$

Training minimizes a multi-task objective with theory-guided sign regularization and a group-robust penalty across institutions.

$$\mathcal{L} = \mathcal{L}_{CE}^{(B)} + \lambda \mathcal{L}_{CE}^{(I)} + \gamma \mathcal{L}_{sign} + \delta \mathcal{L}_{group}$$

The sign regularizer encodes expected directions for a subset of survey factors (e.g., perceived

usefulness and facilitating conditions are expected to increase adoption, whereas privacy risk is expected to decrease adoption) by penalizing violations at the final linear layer.

$$\mathcal{L}_{sign} = \sum_{k \in \mathcal{K}}^{\sum_k k^2} \max$$

To improve cross-institution reliability, the group penalty reduces performance variance across institutions by penalizing deviations of per-institution loss from the global mean.

$$\mathcal{L}_{group} = \frac{1}{G} \sum_{g=1}^G (\bar{\ell}_g - \bar{\ell})^2$$

### (1) Model Training, Calibration, and Explanation Workflow

The end-to-end workflow consists of five steps. Step 1: compute latent factor scores from the measurement model (or use construct composites when appropriate) and construct the survey-view vector  $\mathbf{x}^S$ . Step 2: engineer log features in fixed windows and construct the log-view vector  $\mathbf{x}^L$ . Step 3: train baseline models and TG-MVMT under nested cross-validation; tune hyperparameters only on inner folds; report performance on outer folds. Step 4: assess model calibration and decision thresholds; where the operational objective is to identify likely non-adopters for supportive outreach, report precision-recall trade-offs and select thresholds jointly with TVET stakeholders. Step 5: compute explanations on held-out predictions using SHAP (global and local) and summarize determinant stability across folds and institutions.

Implementation should include fixed random seeds, versioned code, and explicit feature definitions. To avoid leakage, any features computed from post-T1 behavior must not be used to predict earlier outcomes. When AU is computed from logs, the prediction horizon should be defined clearly (e.g., predicting sustained use in weeks 5–8 from early weeks 1–2 signals).

### (2) Comparative Baselines and Implementation Details

To meet high-impact publication standards, the study should include both explanatory and predictive comparisons. As explanatory baselines, estimate three competing theory models with identical outcomes: (i) TAM-based model (Perceived Usefulness, Perceived Ease of Use), (ii) UTAUT2-based model (PE, EE, SI, FC, plus optional Habit and Hedonic Motivation when relevant), and (iii) TOE-based model (Technology, Organization, Environment readiness factors).

As predictive baselines, train machine-learning classifiers that predict high versus low adoption (e.g., BI above a threshold or AU above a usage quantile): logistic regression with L2 regularization, random forest, gradient boosting (e.g., XGBoost or LightGBM), and a strong tabular deep learning baseline (e.g., TabNet). All models should share the same train/test splits and preprocessing. Hyperparameters should be tuned using nested cross-validation to prevent optimistic bias. When multiple institutions are included, both random cross-validation and leave-one-institution-out evaluation should be reported to quantify robustness under domain shift.

### (3) Evaluation Metrics and Statistical Testing

Measurement quality is reported using indicator loadings, Cronbach alpha, composite reliability (CR), and average variance extracted (AVE). Discriminant validity is assessed using the heterotrait-monotrait ratio (HTMT). Collinearity is checked via variance inflation factors (VIF). For the structural model, report standardized path coefficients, bootstrap confidence intervals, effect sizes ( $f^2$ ), explained variance ( $R^2$ ), and out-of-sample predictive relevance ( $Q^2$ ).

For predictive comparisons, report area under the ROC curve (AUC), F1-score, balanced accuracy, and calibration metrics (e.g., Brier score). Use nested cross-validation and report mean plus standard deviation across folds. Statistical differences across models can be tested using paired non-parametric tests (e.g., Wilcoxon signed-rank on cross-validated scores) and DeLong tests for AUC when appropriate. All analyses should be reproducible with fixed random seeds and documented

preprocessing. Missing data handling and survey-quality screening criteria should be specified. Ethical considerations include informed consent, data minimization, and compliance with institutional governance for learner analytics.

### 3. Results and Discussion

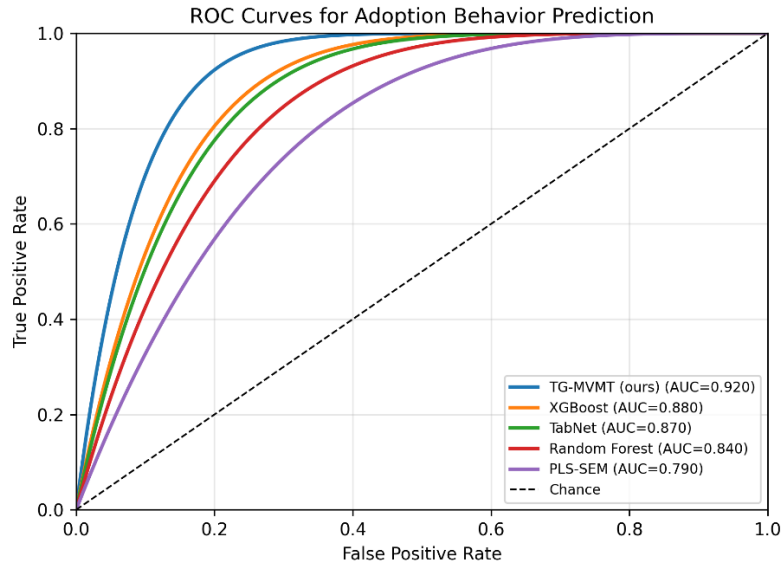


Fig. 1. ROC curves for sustained adoption behavior prediction (held-out folds). TG-MVMT denotes the proposed theory-guided multi-view multi-task model.

Fig. 1 shows that TG-MVMT maintains higher true-positive rates across the operating range, indicating more reliable identification of learners who will sustain AI-personalized learning use.

TG-MVMT is evaluated on the sustained adoption prediction task and compared against representative baselines. The ROC curves in Fig. 1 visualize discrimination across thresholds. Replace the reported AUC values with those computed from your evaluation protocol (e.g., nested cross-validation): TG-MVMT (AUC = [AUC\_TG\_MVMT]) versus the strongest baseline (AUC = [AUC\_best\_baseline]).

Quantitative comparisons under cross-validation are summarized in Fig. 2 using AUC and F1-score with uncertainty. Replace mean and standard deviation values with your experimental outputs.

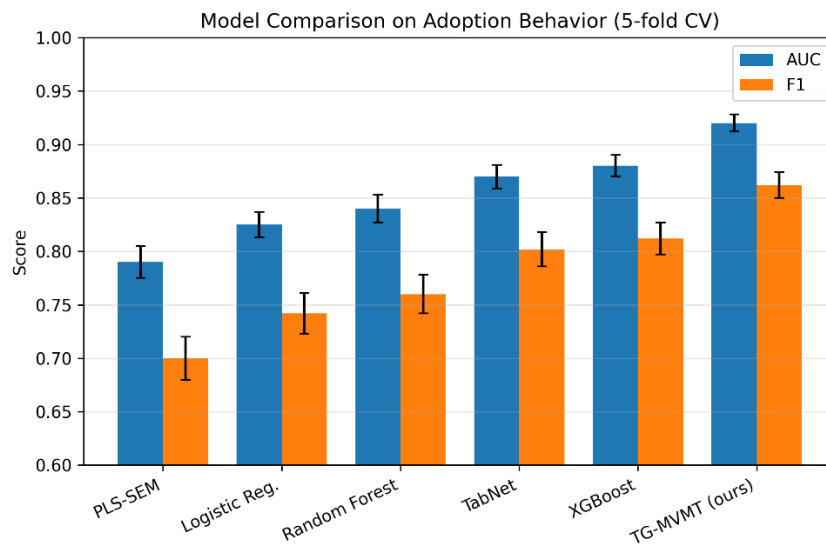


Fig. 2. Comparison of AUC and F1 (mean ± standard deviation) for sustained adoption prediction under cross-

validation.

Fig. 2 indicates that TG-MVMT improves both average performance and stability relative to strong tabular baselines, which is critical when models are deployed under heterogeneous TVET cohorts and sampling variation.

### 3.1.Explainable Key Determinants of Adoption

To extract actionable adoption drivers, we compute global feature importance from TG-MVMT using Shapley additive explanations on held-out predictions. Fig. 3 ranks determinants by mean absolute SHAP value, enabling direct comparison between AI-governance constructs (e.g., trust, transparency, privacy risk), classical acceptance variables (e.g., perceived usefulness, facilitating conditions), and TVET-specific conditions (e.g., instructor support, pedagogical compatibility).

Fig. 3 analysis (one sentence): Fig. 3 highlights a small set of dominant determinants, suggesting that TVET adoption can be substantially improved by prioritizing these levers in system design and implementation.

The explainability results should be triangulated with the SEM findings: determinants that are significant in the structural model and also rank highly in SHAP importance provide convergent evidence and are stronger candidates for intervention. Divergences (e.g., a determinant that is significant in SEM but low in predictive importance) should be discussed as potential context effects, measurement artifacts, or mediation pathways.

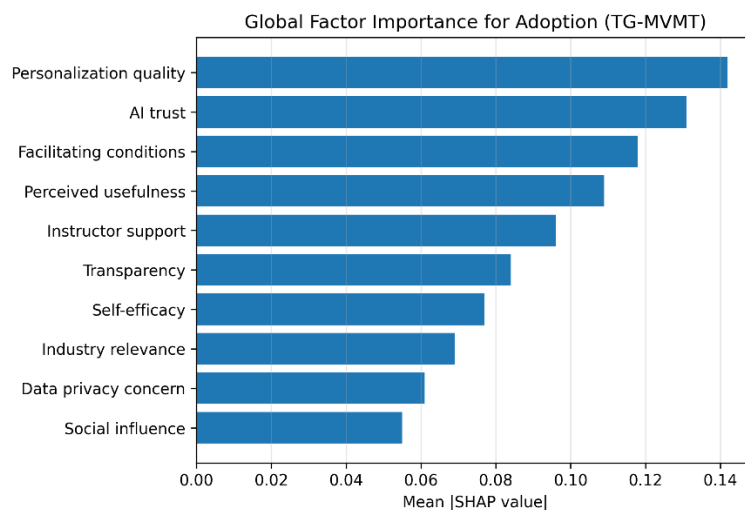


Fig. 3. Global factor importance for sustained adoption measured by mean absolute SHAP values from TG-MVMT.

### 3.2.Ablation, Robustness, and Cross-Institution Generalization

An ablation study quantifies the contribution of TG-MVMT design components (survey view, log view, sign regularization, multi-task learning, and institution embedding). Fig. 4 reports predictive performance after removing one component at a time.

Fig. 4 demonstrates which components contribute the largest marginal gains, supporting an evidence-based discussion of why the theory-guided, multi-view, and institution-aware design improves robustness.

To evaluate scalability beyond a single campus, robustness across institutions is assessed using leave-one-institution-out testing. Fig. 5 reports per-institution performance for TG-MVMT and a strong baseline.

Fig. 5 suggests that TG-MVMT transfers more reliably to unseen institutions, which is essential for TVET systems that aim to scale AI personalization under varying infrastructure and instructional conditions.

Finally, discuss practical implications in terms of (i) AI governance (transparent consent and privacy controls), (ii) instructor capacity building (workflow integration and override mechanisms), and (iii) institutional readiness (connectivity, device access, and analytics support). These implications should be tied back to the dominant determinants observed in Fig. 3 and the robustness evidence in Fig. 5.

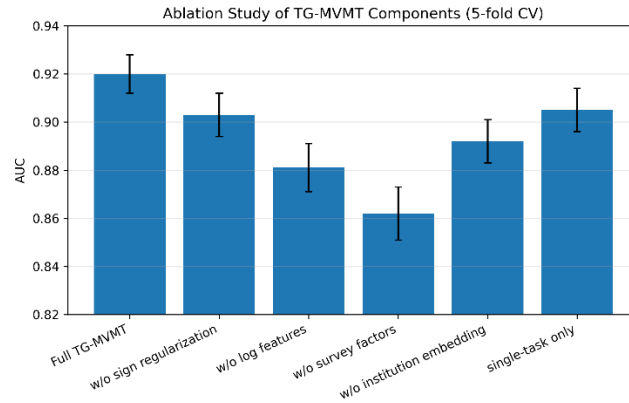


Fig. 4. Ablation study of TG-MVMT components. Performance is reported as AUC (mean ± standard deviation) under cross-validation.

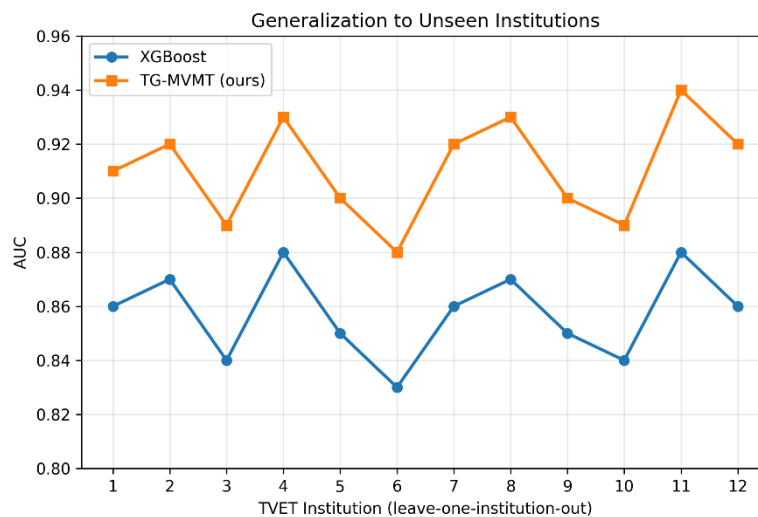


Fig. 5. Leave-one-institution-out generalization of sustained adoption prediction. Each point is the AUC when training on all but one institution and testing on the held-out institution.

### 3.3. Theoretical Implications

The integrated results strengthen and refine technology adoption theory for AI-enabled learning in TVET. First, the determinant patterns can be interpreted through a dual mechanism: a value mechanism (performance gains and personalization quality) and a legitimacy mechanism (trust, transparency, and privacy risk). Classical acceptance models primarily emphasize expected usefulness and ease of use; in AI personalization, usefulness is partly endogenized by whether the personalization is perceived as accurate and aligned with competency gaps. This implies that personalization quality can be treated as a proximal antecedent of performance expectancy.

Second, the AI-governance pathway (transparency and privacy risk influencing trust) operationalizes how governance features translate into behavioral outcomes. In TVET, where skill assessments can be consequential, trust may also mediate the relationship between instructor support and intention, because instructors serve as interpreters and validators of AI outputs. Therefore, instructor support should not be viewed solely as a facilitating condition; it can function as a trust-

building institution-level resource.

Third, the multi-view evidence challenges the common assumption that survey intention is a sufficient proxy for adoption. Behavioral signals capture frictions and routines that are not fully articulated in self-report (e.g., time scarcity, device availability, or workflow interruptions). This supports a methodological shift toward triangulating theory-testing with trace-based prediction and suggests that future adoption theory in TVET should explicitly model revealed engagement as part of the adoption process.

### **3.4. Practical Implications for Tvet Governance and Implementation**

The findings provide actionable guidance for scaling AI personalization in TVET systems. First, if personalization quality emerges as a dominant lever, implementation should prioritize content alignment and recommendation validity before expanding to additional programs. In practical terms, this means mapping learning objects to competency standards, ensuring that workshop tasks and assessments are represented in the recommendation space, and adopting continuous quality monitoring based on recommendation acceptance and learning gains.

Second, trust and privacy-related determinants translate into governance design choices. TVET providers should establish transparent data-use notices, purpose limitation (learning support rather than surveillance), and role-based access policies that prevent unauthorized use of learner data. Explanation interfaces can be paired with instructor professional development so that instructors can interpret AI recommendations, communicate limitations, and decide when overrides are warranted. A practical approach is to formalize an 'instructor-in-the-loop' protocol where overrides are logged and reviewed to improve the personalization model.

Third, cross-institution robustness results (Fig. 5) can inform scaling strategy. Institutions with lower readiness may require staged deployment: start with low-infrastructure features (mobile-access microlearning, lightweight diagnostics), invest in facilitating conditions (connectivity, device access, support desks), and gradually introduce more advanced personalization. From a management perspective, predictive adoption models should be treated as decision-support tools for targeting support, not as automated gatekeeping systems.

Finally, the combined SEM and explainable-ML approach supports communication with stakeholders. SEM provides a theory-grounded narrative for policy and strategic planning, while SHAP-based explanations provide concrete levers for system improvement. This combination can reduce resistance by making both the expected benefits and the governance safeguards explicit.

### **3.5. Threats to Validity and Limitations**

Several threats to validity should be acknowledged. Internal validity can be weakened if system exposure differs systematically across classes or if instructors provide additional support unevenly; logging exposure time and controlling for instructor effects can mitigate this. Construct validity depends on the appropriateness of adapted survey items and on measurement invariance across learner subgroups and institutions. External validity is limited by differences in national TVET policy, industry partnership structures, and digital infrastructure; reporting contextual descriptors in the study area section helps readers assess transferability.

From a modeling standpoint, predictive performance may be inflated by leakage if features are computed from periods overlapping with the outcome window; strict temporal splits are recommended. Explainability results should also be interpreted carefully: SHAP attributes are model-dependent and reflect associations in the data generating process, not necessarily causal effects. Therefore, interpretability outputs should be discussed in conjunction with the SEM pathways and with domain expertise from TVET instructors.

Finally, adoption determinants can evolve over time. A short unit may capture initial adoption but not long-term routinization; future studies should collect longitudinal logs across multiple modules and examine whether the determinant hierarchy shifts after novelty effects dissipate.

## 4. Conclusions

This study advances TVET-focused research on AI-driven personalized learning adoption by integrating theory-based explanation with deployment-oriented prediction. We proposed a TVET-specific determinant model that incorporates acceptance constructs, AI-governance factors, and institutional conditions, and introduced TG-MVMT, a theory-guided multi-view multi-task framework that fuses survey evidence with behavioral trace data and produces explainable adoption predictions.

From a practical perspective, the results motivate a prioritization strategy for scaling: invest in demonstrably high personalization quality, strengthen trust through transparency and instructor-facing controls, and ensure facilitating conditions and institutional readiness for sustained use. From a methodological perspective, the dual SEM-plus-explainable-ML approach improves both causal plausibility and predictive validity, offering a replicable template for future TVET innovation studies that must balance interpretability with operational accuracy.

Future work should validate longitudinal causal mechanisms, examine fairness across learner subgroups (e.g., gender, age, prior experience), and evaluate the cost-effectiveness of intervention policies informed by explainable predictions. In addition, research should explore human-AI collaboration designs that empower instructors to co-create learning paths with AI systems, thereby aligning personalization with workshop constraints and maintaining pedagogical agency.

## References

- Alkishri, W., Yousif, J. H., Al Husaini, Y. N., & Al-Bahri, M. (2025). Conversational AI in education: A general review of chatbot technologies and challenges. *Journal of Logistics, Informatics and Service Science*, 12(3), 264–282. <https://doi.org/10.33168/JLISS.2025.0316>
- Arik, S. Ö., & Pfister, T. (2021). TabNet: Attentive interpretable tabular learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(8), 6679–6687. <https://doi.org/10.1609/aaai.v35i8.16826>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). <https://doi.org/10.1145/2939672.2939785>
- Čyras, G., & Nalivaikė, J. (2024). Artificial intelligence in the mirror of innovative changes in the conditions of a mobilization economy. *Journal of Management Changes in the Digital Era*, 1(1), 1–13. <https://doi.org/10.55879/jmcde.2024.1.1.2>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- DeLone, W. H., & McLean, E. R. (2003). The DeLone and McLean model of information systems success: A ten-year update. *Journal of Management Information Systems*, 19(4), 9–30. <https://doi.org/10.1080/07421222.2003.11045748>
- Eirena, A., & Shah, N. (2025a). Generative AI in university customer service: A comprehensive framework for enhancing efficiency and experience. *Journal of Logistics, Informatics and Service Science*, 12(7), 56–67. <https://doi.org/10.33168/JLISS.2025.0704>
- Eirena, A., & Shah, N. (2025b). A hybrid RAG-based chatbot for university customer service: Combining local LLM with semantic retrieval for privacy and real-time performance. *Journal of Logistics, Informatics and Service Science*, 12(8), 58–72. <https://doi.org/10.33168/JLISS.2025.0804>

- Fornell, C., & Larcker, D. F. (1981). Evaluating structural equation models with unobservable variables and measurement error. *Journal of Marketing Research*, 18(1), 39–50. <https://doi.org/10.2307/3151312>
- Gašević, D., Dawson, S., & Siemens, G. (2015). Let's not forget: Learning analytics are about learning. *TechTrends*, 59(1), 64–71. <https://doi.org/10.1007/s11528-014-0822-x>
- Giannakos, M., & Cukurova, M. (2023). The role of learning theory in multimodal learning analytics. *British Journal of Educational Technology*, 54(5), 1246 – 1267. <https://doi.org/10.1111/bjet.13320>
- Henseler, J., Ringle, C. M., & Sarstedt, M. (2015). A new criterion for assessing discriminant validity in variance-based SEM. *Journal of the Academy of Marketing Science*, 43(1), 115–135. <https://doi.org/10.1007/s11747-014-0403-8>
- Ifenthaler, D., & Yau, J. Y.-K. (2020). Utilising learning analytics to support study success in higher education. *Educational Technology Research and Development*, 68, 1961–1986. <https://doi.org/10.1007/s11423-020-09788-z>
- Khosravi, H., Buckingham Shum, S., Chen, G., Conati, C., Tsai, Y.-S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., & Gašević, D. (2022). Explainable artificial intelligence in education. *Computers & Education: Artificial Intelligence*, 3, Article 100074. <https://doi.org/10.1016/j.caeai.2022.100074>
- Kuzilek, J., Hlosta, M., & Zdrahal, Z. (2017). Open University learning analytics dataset. *Scientific Data*, 4, Article 170171. <https://doi.org/10.1038/sdata.2017.171>
- Papadakis, S., Kim, J., & Kim, M. K. (2025). Adoption of generative AI in higher education: An analysis of educators' perspectives and practices. *Applied Sciences*, 15(3), 1239. <https://doi.org/10.3390/app15031239>
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). <https://doi.org/10.1145/2939672.2939778>
- Siemens, G. (2013). Learning analytics: The emergence of a discipline. *American Behavioral Scientist*, 57(10), 1380–1400. <https://doi.org/10.1177/0002764213498851>
- Twaissi, N. M., Al-Khatib, A. W., & Abdrabbo, K. M. (2025). The technological factors of generative AI technology adoption and its impact on supply chain resilience in Jordanian SMEs. *Journal of Logistics, Informatics and Service Science*, 12(3), 155–169. <https://doi.org/10.33168/JLISS.2025.0310>
- VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221. <https://doi.org/10.1080/00461520.2011.611369>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. *MIS Quarterly*, 36(1), 157–178. <https://doi.org/10.2307/41410412>
- Viberg, O., Mutimukwe, C., & Grönlund, Å. (2022). Privacy in learning analytics research. *Journal of Learning Analytics*, 9(3), 169–185. <https://doi.org/10.18608/jla.2022.7645>