

Personalised Learning Pathways for Students Using Adaptive AI and Limited Educational Data

Sanchita Kiran

Department of Computer Science and Engineering, Silicon University, Bhubaneswar, Odisha, India

sanchitakiran08@gmail.com

Abstract. Adaptive systems based on Artificial Intelligence (AI) are highly dependent on the paradigm of personalised education, which focuses on tailoring learning experiences to the individual needs of students. The major obstacle to effective and widespread adoption of these systems is that they require high volumes of student interaction data. This reliance poses a critical issue in data-sparse settings, as with new students (the cold-start problem), on niche or advanced courses with limited enrolment, or where data privacy rules are highly restrictive. As a result, the possibility of real personalisation is not realised in many cases of practical education. To overcome this inherent setback, a new Hybrid Meta-Transfer Learning (HMTL) framework is presented. The HTML framework is designed to create effective, customised learning pathways with minimal initial data. It does so by a synergistic, two-stage architecture: a Transfer Learning component, which is used to construct a robust, generalised Foundational Knowledge Core with data-rich source domains, and a Meta-Learning-based Rapid Adaptation Engine, which is then used to specialise this Core to a new student using only a small number of initial interactions. An extensive simulation shows that the suggested framework outperforms baseline models by a very large margin in making predictions about student performance as well as in terms of the efficiency of the produced learning pathways. This paper introduces the HTML framework with its architecture, mathematical formulation, and empirical validation, showing a major step forward in the development of a really personalised learning experience using minimal educational data.

Keywords: Personalised Learning; Adaptive AI; Limited Data; Data Sparsity; Meta-Learning; Transfer Learning; Student Modelling; Educational Data Mining.

1. INTRODUCTION

Modern education is experiencing a radical change, as the focus is moving from the traditional, one-size-fits-all models of instruction to the more dynamic and student-centred model. The central point of such an evolution is the idea of personalised learning, which is the instructional method that varies the pace, material, and format of the learning process according to the needs, advantages, and disadvantages of individual learners. This personalised practice is no longer viewed as a pedagogical luxury but is being acknowledged as an urgent need to increase student engagement, improve their knowledge retention, and eventually improve their academic performance across diverse populations of learners. Artificial Intelligence (AI) has greatly heightened the promise of providing truly personalised education at scale. AI-based adaptive learning systems can gather and process large volumes of multidimensional data, such as performance data in the form of assessments, behavioural data like time-on-task and engagement data, as well as predictive analytics to detect students potentially falling behind in terms of learning, hence allowing timely and efficient interventions. Despite its immense potential, the efficacy of the dominant paradigm in AI-driven education is fundamentally constrained by a critical bottleneck [1]: a profound dependency on large-scale datasets. The sophisticated algorithms that power these adaptive systems typically require extensive histories of student interactions to build accurate predictive models and make reliable pedagogical decisions. This "big data" assumption creates a significant and pervasive challenge in numerous real-world educational contexts, a phenomenon known as data sparsity. Data sparsity manifests in several critical forms. The most prominent is the student cold-start problem, where a new student with no prior interaction history cannot receive any meaningful personalisation from the system. Similarly, niche, advanced, or newly introduced courses often lack the high enrollment numbers necessary to generate a sufficient volume of data for robust model training. Furthermore, the growing importance of data privacy and regulations like FERPA and GDPR rightly impose limits on the scope of data collection, further constraining the amount of information available for model training.

The technical consequences of sparse data are extreme. Under these circumstances, machine learning models become too vulnerable to overfitting and learn the noise in the limited dataset instead of the real patterns of learning, and thus generalise poorly to new interactions. Adaptive strategies based on the multi-armed bandit model, including Thompson Sampling, have additionally been demonstrated to perform worse, resulting in slow convergence and not reliably differentiating between effective and ineffective learning materials in the case of sparse feedback. It is this incompatible mismatch between the data-hungry specifications of present-day AI and the frequently data-sparse nature of the educational context that the promise of personalisation is not achieved by

a significant proportion of learners. The field requires a paradigm shift in favour of data-efficient AI systems as opposed to data-intensive models tailored specifically to the needs of the educational space.

To seal this vital gap, the proposed work presents a novel Hybrid Meta-Transfer Learning (HMTL) framework. The HMTL structure is designed with the fundamental objective of tackling the issue of data sparsity and offering useful personalisation with a small number of interactions with the student. It uses a two-pronged approach that integrates the virtues of two powerful machine learning paradigms. Initially, the Foundational Knowledge Core [5] is developed using Transfer Learning, which essentially trains a model on a large, data-intensive source domain (e.g., a large introductory course with a high enrolment). This enables the model to be trained to give generalised knowledge structures and learning processes. Second, a Rapid Adaptation Engine is developed by means of Meta-Learning or “learning to learn”. This engine uses the powerful, pre-trained foundation model and learns an optimisation process that can fine-tune it to the unique learning patterns of a new student using just a few of their early interactions. The HMTL framework has the potential to produce effective and robust personalised learning pathways, even in the most difficult cold-start conditions, through broad-based generalisation of transfer learning and fast, few-shot adaptation via meta-learning as per Figure 1.




Figure 1. Overview

The three key contributions of the given work the following: 1) the design and formalization of the new HMTL architecture, a data-efficient adaptive learning framework; 2) the empirical validation of the superiority of the framework in predictive accuracy and learning pathway efficiency over the existing baseline models in data-limited predictive learning; 3) the provision of a practical and robust solution to the long-standing cold-start issue of adaptive learning systems. The rest of the paper is designed in the following manner: Section 2 will be a review of the literature regarding adaptive learning systems. Section 3 formally outlines the problems and research goals. The HMTL methodology is outlined in Section 4. The experimental results are given and explained in Section 5. Lastly, Section 6 is a conclusion of the paper and provides future work directions. Figure 1 gives a graphical abstract of the paper.

2. LITERATURE REVIEW

The history of pursuing personalised education has been a long journey, starting with primitive, rule-based Intelligent Tutoring Systems (ITSs) and developing to the data-driven, highly engineered AI of the recent past. The early ITSs relied on expert-driven rules and cognitive models to steer students through a curriculum. These systems were unbending and hard to expand, yet innovative. The modern paradigm makes use of machine learning to design adaptive learning contexts, which are dynamic and receptive to the interactions between individual learners. These modern systems are usually developed based on a collection of fundamental elements that collaborate to provide a personal experience.

Core Components of Modern Adaptive Systems

A modern adaptive learning is a closed-loop of evaluation, modelling, and intervention. Student modelling, content adaptation, and feedback mechanisms are the main elements that contribute to the working of this cycle.

- **Student Modelling:** The student model is a computational depiction of the knowledge, skills, and other attributes of a learner, the backbone of any adaptive system. Knowledge tracing is one of the main objectives of student modelling, which is a problem of tracing how a student manages to master various concepts over time, depending on the history of interaction between the student and the teacher. Early and powerful methods involve Bayesian Knowledge Tracing (BKT) [7] that employs a Hidden Markov Model to model the knowledge state of a student per-concept as a latent binary variable (mastered or not mastered). BKT is explainable but tends to over-simplify the learning process by supposing that the skills are independent. More recent enhancements have used deep learning, and in particular Deep Knowledge Tracing (DKT) [8], which uses Recurrent Neural Networks (RNNs) to predict the sequence of interactions between a student and a teacher. DKT is capable of learning detailed, non-linear relations between abilities, without any feature engineering by hand, yet it is infamously data-intensive and overfits in sparse information contexts, which is a major drawback to its practical implementation [9-11].
- **Content Adaptation:** Based on the current state of the student model, the system will be able to dynamically choose and deliver the best learning material. This content-adaptation process may include scaling the difficulty of the next problem, suggesting a remedial video, or proposing a more challenging topic. Predictive analytics are frequently used to estimate how well a student is likely to perform on a series of candidate items, so the system can select the item that maximises a particular pedagogical goal, such as anticipated knowledge acquisition or engagement. One such common framework used to optimise this sequential decision-making process is the Multi-Armed Bandit (MAB) [2-4], which offers a principled approach to trade-off between exploiting already known to be effective content and exploring new content to find potentially better alternatives.
- **Feedback Mechanisms:** The last activity in the adaptive loop is the input of feedback. AI-driven systems may provide feedback on student responses on a granular scale in real time, a capability that is essential to effective learning but often impractical for human instructors to provide at a large scale. Such feedback may be as straightforward as a correct/incorrect indication, or elaborate instructions or hints produced by the natural language processing models [12]. The loop of constant feedback enables quick refinement of the student model and directs future decisions of content adaptation.

The Pervasive Challenge of Data Sparsity in Educational Data Mining

Although the above-described components are highly effective as a theoretical framework that can be used in personalisation, they are severely lacking in performance in practice due to the problem of data sparsity, which is prevalent. Within the framework of Educational Data Mining (EDM), sparsity is the large percentage of missing data in the student-interaction matrix [13-15]. An example is when a matrix is used with the rows depicting students and the columns depicting learning concepts or items. This is because there is an entry in the cell only when a student has interacted with a particular item. A student in any realistic course will, in fact, touch only a tiny fraction of the overall available items, so that most of the matrix is vacant. Real-world data of learning performance may include up to 80-90 per cent of missing measurements. The result of this sparsity is extreme technical difficulties [16]. It further increases the “curse of dimensionality,” in which the set of potential learning pathways or sequence of content increases exponentially, and the data to estimate this is limited. Having small samples, it is statistically hard, even impossible, to reliably differentiate the effectiveness of various pedagogical strategies or even content variations [17-18].

This has a direct effect on the optimisation algorithms, such as MABs, either failing to reach an optimal policy or arriving at a suboptimal policy because of the noisy and sparse reward signals. The most severe occurrence of this problem is the cold-start problem, in which the system initially has no data concerning a new student or a new course, making data-based personalisation impossible on the first try. This contradiction is not merely a little inconvenience but a deficiency that causes adaptive systems not to be effective at the time when they are needed most of all: when a learner is starting their learning path. Table 1 gives a comparative study of the leading technologies applied in adaptive learning with references to their concepts and, most importantly, limitations in the presence of the data sparsity issue. It is identified that each paradigm alone cannot solve this issue, which means that a hybrid approach should be developed, which is capable of integrating the merits of the various techniques into a more robust and data-efficient solution, as per Table 1.

Table 1: Comparative Analysis of Technologies for Adaptive Learning

Technology/Paradigm	Core Principle	Key Features for Personalisation	Limitations in Data-Sparse Contexts
Bayesian Knowledge Tracing (BKT) [7]	Models student knowledge as a latent binary variable (mastered/not mastered) using a Hidden Markov Model.	Tracks mastery of individual knowledge components; interpretable parameters (slip, guess).	Assumes independence between skills; struggles with complex, multi-skill concepts; requires sufficient data per skill.
Deep Knowledge Tracing (DKT) [8]	Uses Recurrent Neural Networks (RNNs) to model the temporal sequence of student interactions.	Captures complex, non-linear relationships and skill dependencies without explicit feature engineering.	Highly data-hungry; prone to overfitting with sparse data; acts as a "black box," lacking interpretability.
Multi-Armed Bandits (MAB) [2]	Balances exploration (trying new content) and exploitation (using known effective content) to optimise a reward signal (e.g., learning gain).	Efficiently identifies optimal content/pathways in real-time.	Standard algorithms (e.g., Thompson Sampling) underperform and converge slowly in sparse-reward/data environments.
Generative AI (e.g., GANs) [12]	Learns the underlying distribution of a dataset to generate new, synthetic data samples.	Can be used for data augmentation to enrich sparse datasets.	Generated data may not perfectly capture the nuances of real student learning; it can be unstable to train; and it introduces its own set of biases.
Standard Supervised Models	Learns a direct mapping from student features to outcomes (e.g., performance prediction).	Can predict at-risk students or final grades with high accuracy, given sufficient data.	Suffers severely from overfitting and poor generalisation when data is limited; cannot adapt dynamically to new students without retraining.

Learning with limited data has been the subject of widespread research in machine learning, with meta-learning, transfer learning, and hybrid deep learning architectures having provided great potential solutions. “Learning to learn” (also known as meta-learning) is concerned with the ability to make models quickly adapt to new tasks using very few examples. Methods such as Model-Agnostic Meta-Learning (MAML) are trained to minimise the model parameters to adapt rapidly. They are thus suitable for personalised learning where a new student represents a new task. The meta-training process can eliminate the cold-start problem by training on data from a variety of students, and the model has demonstrated potential in similar directions, like personalised federated learning. In parallel, transfer learning tackles data sparsity by pre-training models on large, data-rich tasks and fine-tuning them for smaller, related ones. In education, this could involve pre-training on widely taught courses (e.g., Algebra I) and adapting the model for specialised or low-enrollment courses, a strategy that has improved knowledge tracing models with sparse interaction data.

Another dimension is hybrid deep learning architectures that combine the advantages of models like CNNs to extract features and RNNs or GRUs to model the sequential information. These synergistic designs have proven to have better performance in the prediction of student performance than single models. Nevertheless, the current literature tends to consider meta-learning and transfer learning as two independent categories: transfer learning produces generalised representations and is not flexible, whereas meta-learning is flexible but highly sensitive to the quality of initialisation. Their combination provides a very attractive synergy, namely to offer a powerful and knowledge-rich base by transfer learning and then adapt it to the needs of individual students through meta-learning. It is this combination that gives the theoretical foundation of the proposed HTML framework, bringing these complementary paradigms to work together in data-efficient personalised learning.

3. PROBLEM STATEMENT & RESEARCH OBJECTIVES

Problem Statement

Given a set of learning resources available, $L = \{l_1, l_2, \dots, l_m\}$ and a new student s who has a very limited interaction history modelled by $H_s = \{(l_j, o_j)\}_{j=1}^k$, where k is small (e.g., $k < 5$) and $o_j \in \{0, 1\}$ is the binary outcome (like incorrect/correct), the problem here is to generate a learning pathway $P_s = \{l_{p1}, l_{p2}, \dots, l_{pn}\}$ That is optimal and personalised. An optimal pathway maximises the expected increase in students. s knowledge and minimises the resources (n) that are required, this optimisation has to be carried out with the very drastic restriction of the data sparsity, in that the initial interaction history is. H_s cannot provide sufficient data to be used by the traditional adaptive algorithms in constructing a reliable student model.

Research Objectives

To address the problem stated above, the following research objectives are established:

1. **To design and formalise a Hybrid Meta-Transfer Learning (HMTL) framework capable of modelling student knowledge and generating effective learning pathways, even when initialised with sparse data.**
2. **To implement the HMTL framework and evaluate its performance** on a simulated educational dataset, comparing its predictive accuracy and pathway efficiency against established baseline models, namely a standard recurrent model and a pure transfer learning model.
3. **To conduct an ablation analysis** to quantify the individual contributions of the meta-learning and transfer learning components to the overall performance of the hybrid framework, thereby validating the synergistic design.
4. **To demonstrate the framework's robustness** by systematically testing its performance across varying levels of data sparsity, with a specific focus on its efficacy in the cold-start scenario ($k=1$).

4. METHODOLOGY

In order to address the problem of creating individualised learning trajectories with scarce information, the Hybrid Meta-Transfer Learning (HMTL) framework is introduced. The HMTL architecture is a two-component framework that aims to synergistically integrate the knowledge generalisation of transfer learning with the fast and few-shot adaptation of meta-learning. The initial part is the Foundational Knowledge Core (FKC), a deep neural network that is pre-trained on a large, data-rich source dataset. This pre-training gives the model strong, generalised knowledge of learning dynamics and the relationships between concepts. The second sub-component is the Rapid Adaptation Engine (RAE) [6], which is based on a meta-learning algorithm to learn an optimisation process that can adjust the FKC effectively to a new, unknown student with just their very limited interaction history.

Mathematical Formulation

The Hybrid Meta-Transfer Learning (HMTL) model is based on a recurrent neural network (RNN), or more precisely, a Gated Recurrent Unit (GRU), that can be used to capture the temporal dynamics of learning among students. What is known by the student at the moment t , denoted as, h_t is expressed by the hidden state of the GRU.

Component 1: Foundational Knowledge Core (Transfer Learning)

The Foundational Knowledge Core (FKC) is initially trained on a large source dataset D_{source} that contains long interaction sequences of many past students. This step solves a set of early model parameters ϕ , which can represent overall patterns of acquisition of knowledge. A particular interaction at time t is modelled as a vector. x_t , which concatenates the embedding of a learning item and the result of the prior interaction. The GRU update equations are given as follows in (1), (2), (3), (4):

$$z_t = \sigma(W_z x_t + U_z h_{t-1} + b_z) \quad (1)$$

$$r_t = \sigma(W_r x_t + U_r h_{t-1} + b_r) \quad (2)$$

$$\underline{h}_t = \tanh(W_h x_t + U_h (r_t \odot h_{t-1}) + b_h) \quad (3)$$

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \underline{h}_t \quad (4)$$

The prediction for the next outcome is given by (5):

$$\widehat{y_{t+1}} = \sigma(W_y h_t + b_y) \quad (5)$$

The binary cross-entropy (BCE) loss over all sequences in D_{source} is shown in (6):

$$L_{BCE}(y, \hat{y}) = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})] \quad (6)$$

The optimal transfer learning parameters are obtained as given in (7):

$$(7) \quad \phi^* = \arg \min_{\sum_{(X,Y) \in D_{source}} \sum_t L_{BCE}(y_t, \hat{y}_t)}$$

Component 2: Rapid Adaptation Engine (Meta-Learning)

The **Rapid Adaptation Engine (RAE)** uses the pre-trained parameters ϕ^* as initialisation with the goal of enabling rapid adaptation to any new student task using a few interactions.

Each task T_i corresponds to one student and is split into:

- A support set $D_{T_i}^{supp}$ That is for adaptation
- A query set $D_{T_i}^{query}$ That is for evaluation.

Following the Model-Agnostic Meta-Learning (MAML) algorithm, for each task, the model is adapted by gradient descent on the support set as shown in (8):

$$\theta'_i = \phi^* - \alpha \nabla_{\phi^*} L_{T_i}(f_{\phi^*}; D_{T_i}^{supp}) \quad (8)$$

The adapted parameters θ'_i are then evaluated on the query set. The meta-objective then updates ϕ^* as shown in (9):

$$(9) \quad \phi^* \leftarrow \phi^* - \beta \nabla_{\phi^*} \sum_{T_i \sim p(T)} L_{T_i}(f_{\phi^*}; D_{T_i}^{query})$$

The update is performed with the help of the Adam optimiser as shown in (10), (11), (12):

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t \quad (10)$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \quad (11)$$

$$\phi_t^* = \phi_{t-1}^* - \eta \frac{\widehat{m}_t}{\sqrt{\widehat{v}_t + \epsilon}} \quad (12)$$

Here g_t is the gradient from the meta-objective.

Pathway Generation at Inference

For a new student who has few initial interactions H_s , the meta-trained parameters ϕ^* are adapted using the inner loop update as given in (13):

$$\theta'_s = \phi^* - \alpha \nabla_{\phi^*} L(f_{\phi^*}; H_s) \quad (13)$$

After this, the system generates a personalised learning pathway iteratively. At each step, it evaluates all candidate learning resources and selects the one that maximises the probability of success (14):

$$l_{next} = \arg \max_{l_j \in L_{candidate}} P(y = 1 | H_s, l_j; \theta'_s) \quad (14)$$

Algorithmic

Frameworks

The logic is summarised in Algorithm 1 and Figure 2, given below:

Algorithm 1: HTML Pathway Generation

Input: New student's initial interactions H_s

Output: Personalised Learning Pathway P

```

1:  $\theta's \leftarrow \text{Adapt}(\phi^*, H_s)$   # few-shot update using Eq. (8)
2:  $P \leftarrow \emptyset$ 
3: while not MasteryAchieved( $P$ ) do
4:   for each candidate resource  $l_j \in L$  do
5:     Compute  $\text{score}_j = P(y=1 \mid H_s, l_j; \theta's)$ 
6:   end for
7:    $l_{next} \leftarrow \text{argmax}(\text{score}_j)$ 
8:   Append  $l_{next}$  to pathway  $P$ 
9:   Update  $H_s$  with student's interaction on  $l_{next}$ 
10:  Update  $\theta's$  via  $\text{Adapt}(\theta's, \text{new interaction})$ 
11: end while
12: return  $P$ 

```




Figure 2. Workflow

5. RESULTS & DISCUSSIONS

The effectiveness of the suggested Hybrid Meta-Transfer Learning (HMTL) framework was validated through a set of controlled experiments conducted on a synthetic dataset. This approach enables an objective assessment of the model's performance at different, specifically controlled data sparsity levels, directly probing its ability to address the cold-start problem. To indicate that it outperforms the two other applicable baseline models in predictive accuracy and learning pathway efficiency, the HTML framework was compared with them. Those findings validate that the synergistic effect of transfer and meta-learning offers significant benefit in data-limited learning.

Experimental Setup

Dataset: A synthetic dataset was created with 50 interdependent concepts and 200 learning items. It included 1,000 "source" students (average 100 interactions each) for pre-training and 200 "target" students (1–10 interactions each) for meta-training and testing, simulating cold-start scenarios. Each student had a unique learning rate, affecting their mastery probability after correct interactions.

Baseline Models:

1. **Standard GRU:** GRU-based network trained from scratch on the available sparse data on a particular target student. This model depicts the performance of a typical deep learning application to the cold-start issue with no prior knowledge.
2. **Transfer Learning Only (TL-GRU):** A GRU model pre-trained on the large source dataset (D_{source}) and then conventionally fine-tuned on the sparse data of the target student. This baseline isolates the contribution of transfer learning alone.

Evaluation Metrics:

1. **Predictive Accuracy (AUC):** The Area Under the Receiver Operating Characteristic Curve for the task of predicting whether a student will answer the next unseen item correctly. This measures the quality of the underlying student model.
2. **Average Learning Gain:** The average increase in the number of mastered concepts for a student after completing the generated pathway.
3. **Average Path Efficiency:** The learning gain divided by the number of items in the generated pathway, measuring how quickly the model can guide a student to mastery.

Quantitative and Comparative Analysis

The overall performance of the HTML framework versus the baselines is exhibited in Table 2. It suggests the HTML model is by far better than the others in all the essential metrics. Standard GRU with access to only sparse data fares poorly, having an AUC value just above chance (0.55), reflecting the failure to construct a useful student model. There is a great improvement in the TL-GRU as the knowledge obtained in the source domain is transferred to the target domain, and the AUC attains 0.72. Nevertheless, the HTML structure performs significantly better than the two, with an AUC of 0.85. This strong predictive power is directly converted to higher predictive power pathways, where HMTL has the largest average learning gain (0.88) and path efficiency (0.12), meaning that it produces shorter and more effective learning sequences, as per Table 2

Table 2: Overall Performance Comparison

Model	Predictive Accuracy (AUC)	Average Learning Gain	Average Path Efficiency
Standard GRU	0.55 ± 0.04	0.20 ± 0.08	0.02 ± 0.01
TL-GRU	0.72 ± 0.03	0.65 ± 0.05	0.07 ± 0.01
HTML (Proposed)	0.85 ± 0.02	0.88 ± 0.03	0.12 ± 0.01

An ablation analysis was carried out to confirm the hypothesis that transfer learning combined with meta-learning is essential. Table 3 indicates that the elimination of either of these components leads to a considerable reduction in performance. Removal of the meta-learning component (which amounts to the TL-GRU baseline) leads to a 15.3% reduction in AUC. More distinctly, the elimination of the transfer learning (i.e., the application of meta-learning on a randomised initiation) leads to a drop in performance by 20.0%. This affirms that the Foundational Knowledge Core provided by transfer learning is a very important requirement for the Rapid Adaptation Engine to operate, as per Table 3

Table 3: Ablation Analysis of HMTL Components

Component Removed	Predictive Accuracy (AUC)	Performance Drop (%)
None (Full HMTL)	0.85	-
Meta-Learning (i.e., TL-GRU)	0.72	15.3%
Transfer Learning (Meta-Learn from scratch)	0.68	20.0%

Visual Results and Discussion

The performance dynamics are further illustrated in the following plots. Figure 3 shows the predictive accuracy of the models as a function of the number of initial interactions available for a new student. The HMTL model starts with a high AUC even with just one data point and plateaus quickly, demonstrating its effectiveness in the true cold-start scenario. In contrast, the baseline models start near random chance and require significantly more data to achieve modest accuracy. Figure 4 displays the meta-loss during the training of the RAE component, showing stable convergence and indicating that the model successfully learned a generalizable adaptation strategy as per Figures 3 and 4.




Figure 3. Predictive Accuracy vs Initial Interactions




Figure 4. Learning Curve during Meta-Training

Direct comparison of the learning gain and path efficiency for the pathways generated by each model is presented in Figures 5 and 6. HMTL reveals more pathways with higher knowledge acquisition (Figure 5) and with fewer items (Figure 6), which proves that it is a better producer of efficient and effective personalised curricula.

Figure 7 visualises sample paths generated by both HMTL and TL-GRU models for the same student profile. The HMTL pathway has a more rational course as it introduces basic required concepts initially and proceeds to more refined subjects. The TL-GRU route is not as straightforward, implying that the items are either too sophisticated or unnecessary, resulting in a longer and less productive learning process. Figure 8 illustrates distributions of pathway lengths where HMTL gives rise to a skewed distribution in shorter, more efficient paths. Lastly, Figure 9 of the confusion matrices in the next-item prediction task demonstrates that HMTL is better in the true positive rate and also in the false positive rate than TL-GRU, which is why it proves to be more accurate in modelling students.

The collective outcomes are good indications that the HMTL framework effectively solves the data sparsity issue. The transfer learning aspect provides an effective inductive bias, ensuring the model does not make blind guesses. The meta-learning component then sharpens this general knowledge, specialising it for the individual learner with remarkable data efficiency.




Figure 5. Comparative Learning Gain




Figure 6. Comparative Path Efficiency




Figure 7(a). HMTL Pathway (Conceptual) & TL-GRU Pathway (Conceptual)




Figure 8. Distribution of Pathway Lengths




Figure 9. Confusion Matrices (HMTL & TL-GRU)

It is this two-fold process that allows HMTL to provide proper predictions and, therefore, produce high-quality custom pathways, even during the initial interaction.

6. CONCLUSION & FUTURE SCOPE

The paper introduces the Hybrid Meta-Transfer Learning (HMTL) model to address data sparsity, a prevalent issue in AI-based personalised education, especially for new learners in cold-start situations. Older models have difficulties in offering meaningful personalisation unless there is a high amount of data, which is a serious limitation to their practical use. The generalizable nature of transfer learning and the fast adaptability of meta-learning are well combined in our proposed HMTL model to build precise models of a student and generate effective learning trajectories with a small amount of data.

Controlled experiments on a synthetic dataset have shown empirically that HMTL is significantly better than traditional recurrent models and pure transfer learning approaches. It had better predictive accuracy (AUC), better learning gain, and pathway efficiency. The analysis of ablation supported the effect of the two components being synergistic, and it was possible to prove that not only the generalised Foundational Knowledge Core but also the Rapid Adaptation Engine are needed to make the framework successful. Even though our results prove that effective personalisation can be achieved without large datasets, this research has its limitations. The experiments were done against a simulated data set, and the performance of the framework remains to be tested with actual educational data.

In future research, this study has a very high potential. We intend to implement the framework in a real-life educational environment and will integrate it with privacy-preserving algorithms, like federated learning, to resolve the issue of data privacy. In addition, we aim to enhance the explainability of the model to the learners and use multi-modal data to create a more holistic depiction of student behaviour. Such developments will bring us closer to an effective, ubiquitous personalised learning experience.

CONFLICT OF INTEREST

The authors declare no conflicts of interest regarding the current research.

AUTHOR CONTRIBUTION

The first author has solely worked and contributed to the entire paper.

REFERENCES

- [1] Song, H., Musabirov, I., Bhattacharjee, A., Durand, A., Franklin, M., Rafferty, A., & Williams, J. J. (2025). *Adaptive Experiments Under Data-Sparse Settings: Applications for Educational Platforms*. arXiv preprint arXiv:2501.03999.
- [2] Lan, A. S., & Baraniuk, R. G. (2017, March). A contextual-bandit approach to personalised learning action selection. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 2432-2436). IEEE.
- [3] Li, L., Chu, W., Langford, J., & Schapire, R. E. (2010). A contextual-bandit approach to personalised news article recommendation, in *Proceedings of the 19th International Conference on World Wide Web* (pp. 661-670).
- [4] Shi, C., Shen, C., & Yang, J. (2021). Federated multi-armed bandits with personalisation. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics* (pp. 2917-2925). PMLR.
- [5] Pan, S. J., & Yang, Q. (2009). A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10), 1345-1359.
- [6] Finn, C., Abbeel, P., & Levine, S. (2017, July). Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning* (pp. 1126-1135). PMLR.
- [7] Corbett, A. T., & Anderson, J. R. (1994). Knowledge tracing: A model of student learning. *User Modelling and User-Adapted Interaction*, 4(4), 253-278.
- [8] Piech, C., Spencer, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L.: Deep knowledge tracing. *arXiv preprint arXiv:1506.05908* (2015). <https://arxiv.org/abs/1506.05908>
- [9] Antoniou, A., Storkey, A., & Edwards, H. (2017). *Data augmentation generative adversarial networks*. arXiv preprint arXiv:1711.04340.
- [10] Arora, S. (2025). Transforming an AI Decision Support System with Knowledge Graphs & CAG. *International Journal on Engineering Artificial Intelligence Management, Decision Support, and Policies*, 2(2), 15-23.
- [11] Foglino, F., Christiano, P., & Abbeel, P. (2019). Unsupervised Curricula for Visual Meta-Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 32.

- [12] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 27.
- [13] Jiang, Y., Luan, S., & Wang, W. Y. (2019). Hierarchical automatic curriculum learning for sparse reward navigation tasks. *Neurocomputing*, 365, 316-326.
- [14] Klink, P., D'Eramo, C., Peters, J., & Pajarinen, J. (2022, June). Constrained optimal transport for curriculum reinforcement learning. In *International Conference on Machine Learning* (pp. 11501-11516). PMLR.
- [15] Lim, J., Kim, J., Kim, S., Park, N., & Yoon, S. (2024). MetaVers: Meta-Learned Versatile Representations for Personalised Federated Learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (pp. 7643-7652).
- [16] Niranjala, S. H., & Goyal, S. B. (2025). Improving Wine Quality Forecasts Using Dynamic Integral Neural Networks and Optimised Against Interference. *International Journal on Smart & Sustainable Intelligent Computing*, 2(1), 52–64.
- [17] Vadisetty, R., & Suyal, H. (2025). Smart Data Augmentation using Generative Adversarial Networks for Rare Oncological Disease Classification. *International Journal on Smart & Sustainable Intelligent Computing*, 2(2), 65-79.
- [18] Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., & Wierstra, D. (2016). Matching networks for one-shot learning. In *Advances in Neural Information Processing Systems*, 29.