International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

EchoTwin: Privacy-Preserving Adaptive Digital Writing Twin

Mr. Kalai Kumar K
Artificial Intelligence and Data Science,
DMI College of Engineering,
Palanchur, Chennai-600123,India.

Mrs. A. Divya
Artificial Intelligence and Data Science,
DMI College of Engineering,
Palanchur, Chennai-600123,India.

Lokesh R G
Artificial Intelligence and Data Science,
DMI College of Engineering,
Palanchur, Chennai-600123,India.

Dhanushiya K
Artificial Intelligence and Data Science,
DMI College of Engineering,
Palanchur, Chennai-600123,India.

Neshali Kannan
Artificial Intelligence and Data Science,
DMI College of Engineering,
Palanchur, Chennai-600123,India.

Nadera Banu
Artificial Intelligence and Data Science,
DMI College of Engineering,
Palanchur, Chennai-600123,India.

Abstract— **The fast development of large language models has made a major contribution to generation of texts automatically, but, nevertheless, current systems cannot provide a real personal generation of texts that can convey personal style. The paper develops a new framework named EchoTwin: Privacy-Preserving Adaptive Digital Writing Twin, that applies dynamically-learned and replicated characteristics of user-specific writing without involving computationally costly fine-tuning of models. The suggested methodology proposes a disentangled representation system to distinguish content and style to allow effective control over style. It uses a contrastive learning-based encoder to create unique user style embeddings, called Digital Twin Signatures. These embeddings are held in a memory-enhanced structure that paves way to effective retrieval and constant updating. Besides, a lightweight, adapter-based conditioning scheme incorporates stylistic aspects in the outputs of base language models. It also uses reinforcement learning based on human feedback (RLHF-lite) to smooth personalization as time goes on. Local data storage and differential privacy are used to ensure privacy. Experimental findings clearly indicate that accuracy, adaptability, and efficiency of personalizing are all enhanced with significant differences compared to the current technologies, making EchoTwin a scalable and safe way to provide next-generation AI writing systems.**

Keywords—Large Language Models, EchoTwin, Digital Writing, Style Embeddings, Reinforcement Learning.

I. NTRODUCTION

The high rate of artificial intelligence development and especially in the field of natural language processing (NLP) has greatly changed how people relate to the workings of machines. Large Language Models (LLM) have proven to be highly successful in creating coherent and context-specific, grammatically correct text in a variety of writing fields including academic text, work-related communication, and fiction. Such systems are now finding their way into common tools, and help users to write emails, reports and technical documents with little effort. Although effective, the existing AI writing systems are more oriented at the production of a generalized content that serves a wide user base instead of trying to match personal preferences of the user. A significant drawback of an AI-based writing assistant is that it can currently not accurately reproduce the personal writing style of a user [1]. Writing style is always personal with variation of factors including vocabulary used, sentence structure, tone and syntactic structure. Although the modern models are capable of producing fluent text, they may fail to reproduce these nuanced elements of style, and the resultant productions are not at all generic and impersonal [2]. Such a disconnection is even more apparent at a professional and academic level when it comes to consistency in writing style which is paramount in preserving identity and credibility.

To overcome this constraint, various methods have been suggested, such as training pre-trained models with user-specific datasets. Fine-tuning poses major challenges including high computation cost, requirement of massive labeled datasets, and face of overfitting although personalization can be enhanced using fine-tuning [3]. Furthermore, retraining frequently would not be applicable in the real world where the preference of the users would change with time. These limitations render personalization via fine-tuning impractical to scalable and dynamic applications. Data privacy is yet another severe issue that is related to individual AI systems. A majority of these personalization methods are based on central data collection, in which the contents created by users are stored and processed to external servers. This creates a serious concern of privacy and security, more so when handling sensitive information like emails, research documents and personal notes [4]. Such regulatory frameworks as GDPR focus on the essence of data minimization and control over its use, and it is crucial to design the systems that will learn based on user data without exposing it to third parties. Moreover, there are no current AI writing systems providing systems of continuous adaptation within the context of real-time user feedback. In real-world applications, people tend to edit AI generated content to make it more compatible with users tastes. Such corrections, however, are hardly effectively employed to enhance future production [5]. Lack of adaptive learning features means that the system makes similar mistakes and will not be able to adapt with the user behavior of writing. This underscores the importance of models which can learn incrementally through the interaction of the users, without necessarily retraining entirely.

Along with personalization and privacy issues, there is no definite distinction between content and style in existing models [6]. The internal representations of most language

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

models confound semantic information with features of writing style, and it is hard to separately control or modify writing style [7]. This complexity inhibits flexibility and reduces the lack of opportunity to use identical personalization in various situations, including being able to write in a formal academic tone and in informal communication. The other significant limitation is failure to change writing style according to the requirement of the context by existing systems. The style of writing is not fixed; a user writes differently, depending on the purpose and audience of the text. An example is that in academic writing, research has a much different tone than in the case of an email or a social media post [8]. The existing AI applications can hardly seem to understand such contextual differing, and the results can be inappropriate in terms of style to the case of their application.

In order to address these limitations, the paper is offering EchoTwin: Privacy-Saving Adaptive Digital Writing Twin, a new framework that will copy and transform individual writing styles to protect and efficiently adapt to it in a safe and secure way. The system proposed puts forward a Digital Twin concept of writing where a stylistic conduct of the user is learnt and updated continuously [9]. In contrast to conventional methods, EchoTwin does not require complete model retraining but rather, it uses a memory-augmented architecture that allows it to store and retrieve user-specific style information in real-time. EchoTwin adds a disentangled style representation algorithm, which divides description from form features and allows fine-tuning of personalization. It also uses a contrastive learning-based style embedding model to create a unique Augmented Twin Signature of each user [10]. This representation is represented in a separate memory module and is utilized to direct the process of text generation via a conditioning process using lightweight adapters. The system also incorporates a feedback system based on reinforcement learning to constantly adapt personalization based on user engagement, and privacy built in to both local data storage and differential privacy practices. In this way, EchoTwin can resolve the shortcomings of current systems by offering scalable, adaptive and privacy friendly personalized text generation.

II. Literature Review

Jing et al. [11] presented a new stylized data-to-text generation system (StyleD2T) to generate coherent and stylistically suited texts with structured data used in e-commerce applications. The model combines logic planning-enriched data embedding and masked style embedding and unbiased stylized text generation to provide semantic coherence and style diversity. Experiments with a real world Taobao dataset showed that the suggested model did better compared to current methods in creating product descriptions of high quality, in the correct context, and in the correct style. The research also touched upon some of the challenges, including logical coherence, unstructured style representation, and there was biased training data addressed with some of the sophisticated deep learning methods. Nevertheless, the model needs massive annotated datasets, and a high level of computer power to be applied, restricting the routine implementation in limited resources settings.

Hajizadeh et al. [12] studied the influence of digital multimodal composing (DMC) and developing multi-literacy in young learners using a narrative inquiry approach. The research involved twins involved in DMC activities including text-, image-, audio-, and video-based narration of stories. The results showed that DMC had a significant positive learning impact on language acquisition, digital literacy, cultural awareness, and learner autonomy. The incorporation of the multimodal aspects increased the motivation of the students and allowed them to gain multiple communication competencies at the same time. Also, the research pointed that DMC facilitates constructivist learning as it encourages learning through doing, as well as fostering creativity and critical thinking. The research however, has a small sample size, which is limited as well as lack of generalizability because the research only concentrates on two participants.

Huang et al. [13] added a review of the utilization of digital twin technology in wildfire monitoring, simulation, and prediction, emphasizing that it could enhance real-time monitoring and decision-making. The proposed study developed a model of a wildfire digital twin (WFDT), the model is developed by combining real-time sensor data, simulation models and artificial intelligence to predict fire behavior and optimize the emergency response strategy. The results propose that digital twin can be used to improve situational awareness, predict fire spread accurately, and allocate resources effectively. Also, the paper illuminated the progress in the field of fire detection methods, computer models, and data integration systems, which showed the increasing relevance of the digital twin in disaster management. The deployment of the wildfire digital twins is however challenged by the lack of availability of real-time data and high level of complexity of the system which may restrict scalability and actual adoption.

Li et al. [14] suggested a new method to enhance personalized text-generation methods by automatically rewrites the prompts, rather than fine-tuning large language models (LLMs). The paper discussed the practical constraint that very frequently LLMs are available as black-box APIs, which implies that immediate engineering is the only possible optimization technique. The suggested framework will use supervised learning and reinforcement learning to train a prompt rewriter that adjusts vital features like summaries and synthesized user context. Experiments in the email field, product reviews, and social media established that prompts that were rewritten outperformed original prompts as well as single method optimizations by far in terms of quality of the generated items. This research also revealed that optimized prompts can be read by humans, and they can give valuable feedback about how to design manual prompts. The method however is a large search space and has to go through numerous cycles of reinforcement learning, which is computationally expensive and time-consuming.

A systematic review of literature of end-to-end Transformer-based models in textual Natural Language Processing (NLP) tasks was the theme of the study by Rahali and Akhloufi [15]. The paper emphasized the superioritys of Transformers to normal models like RNNs and CNNs because they have a self-attention mechanism that is effective in capturing long-

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

range effects and facilitates parallel execution. The paper classified the different types of Transformer architectures, such as encoder based, decoder based, as well as hybrid models and how they are used to answer diverse tasks like text classification, translation, summarization, and question answering. It also highlighted the significance of pre-training and fine-tuning architectures, which can enable models such as BERT and GPT to perform highly with a limited amount of labeled data. Moreover, the research has examined various attention processes, tokenization, and training goals which lead to better NLP performance. But Transformer-based models need both large datasets and high computing power, which may be a bottleneck to their efficiency and usability in practice.

Based on the literature provided above, it can be concluded that multiple researchers studied the new methods of advanced approaches in these domains: digital multimodal learning, stylized text generation, Transformer-based NLP models, prompt engineering, and digital twin systems. These investigations have been very useful in enhancing performance, automating, and forecasting with modern tools such as deep learning and artificial intelligence, as well as data driven-based frameworks. Nevertheless, most of these methods are either domain-based, computationally bottleneck, required large datasets or were not designed to be adaptable and scalable in real time. In addition, they do not even partially incorporate intelligent, efficient and user-friendly solutions that are able to generalize in various real world situations. Hence, in spite of the significant advances, the available literature has not yet fully covered the practical issues and needs of our proposed system, which has an opportunity gap to a more streamlined, scalable, and efficient one.

III. Materials

In order to build and test the proposed EchoTwin framework, a hybrid set of textual data sources that are publicly available and writing samples posted by users were used to achieve the desired result of simulating a real scenario of personalization. In particular, benchmark datasets like the Enron Email Dataset and the Reddit Comment Dataset were used to obtain a wide variety of writing style in both formal and informal spheres of communication. Moreover, to reflect structured and formal writing patterns, subsets of the Wikipedia Corpus were also added. To simulate personalization in the real-life, synthetic user profiles were developed by clustering texts of specific authors within these sets, in such a way that the system could learn consistent style features. In addition to these publicly accessible datasets, document, editing, and text correction entries of the users stored locally were gradually gathered by the suggested data acquisition module, which guaranteed privacy-considering learning. The hybrid data approach enabled the model to effectively train in generalized linguistics frameworks and, at the same time, to adapt to the given linguistics characteristics of the user, which is a balance between scalability, personalization, and data security.

IV. Methodology

The suggested EchoTwin: Privacy-Preserving Adaptive Digital Writing Twin framework is executed to be based on a highly integrated pipeline that facilitates the effortless conversion of the raw textual information into highly individualized outputs. The methodology goes past the preprocessing stage to adaptive learning using mathematical formulations and architectural details to increase the clarity and reproducibility. The stages are all interlinked in a meticulously planned way so that the refined output of one module is the expertly input to the other so that a smooth and effective flow of data is carried out throughout the network.

A. Data Preprocessing and Representation Pipeline

Initially, the collected raw textual data is subjected to a multi-stage preprocessing pipeline to ensure uniformity and quality. The input data $D = \{d_1, d_2, \ldots, d_n\}$consists of heterogeneous sources such as emails, documents, and user edits. These sources often contain noise, inconsistencies, and formatting variations that must be standardized before further processing.

Raw textual data is first hierarchically processed in a series of steps to maintain uniformity and quality. The input data $D = \{d_1, d_2, \ldots, d_n\}$ is a heterogeneous source, that is, it can be emails, documents, and user edits. Such sources are usually noisy, irregular and might vary in formatting, which has to be standardized first before any further eventual processing. The preprocessing converts the raw input to structured sequences:

The preprocessing transforms raw input into structured sequences:

$$d_i \rightarrow S_i = \{w_1, w_2, \ldots, w_m\}$$

where S_i is fueling tokenized sequences of text representations that reflect the linguistic structure of the text. This stage includes:

- Subword-based tokenization (e.g., BPE/WordPiece), allowing the treatment of rare and unknown words in an efficient manner.
- Sentence boundary identification in order to maintain logical partitions.
- POS tagging and dependency parsing to form grammatical associations.

Maintenance of the stylistic features, including punctuation, capitalization and spacing which are very vital to learning the style.

This is represented by each token being encoded as a dense vector:

$$x_i = \text{Embedding}(w_i) \in \mathbb{R}^d \quad (1)$$

where $d = 768$ rep is the dimension of semantic encoding of context.

The obtained processed sequences are added to the incremental buffer allowing a continuous update:

$$B_t = B_{t-1} \cup \{S_i\} \quad (2)$$

The buffer is vital as it ensures continuity in time; hence, the system requires enhancement as new user input is introduced without affecting the past trends.

B. Style Feature Extraction and Disentangled Representation

The resulting processed embeddings are fed into a two-branded neural network that disentangles content and style, so that personalization modulates only stylistic elements but not meaning.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

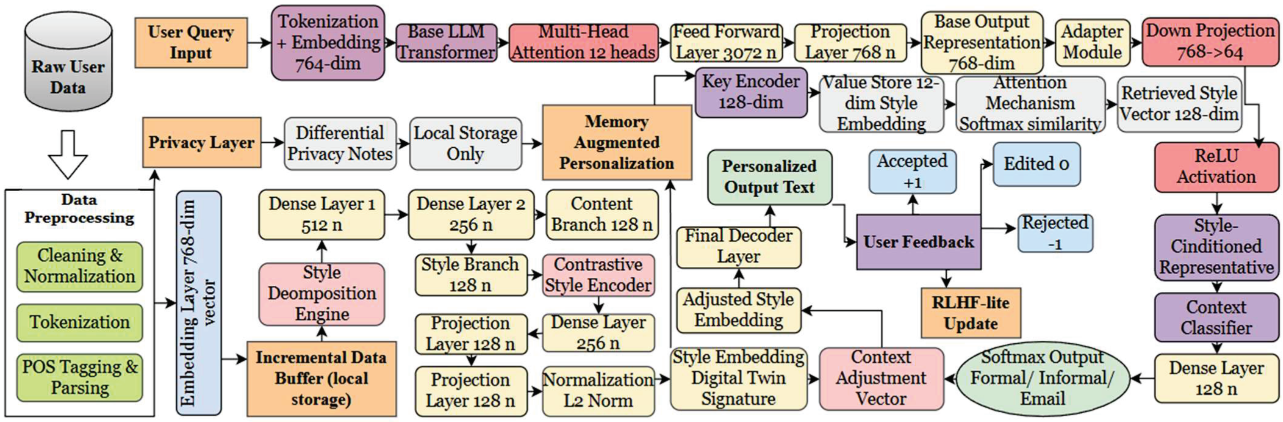

Figure 1: Architecture of the Proposed EchoTwin: Privacy-Preserving Adaptive Digital Writing Twin Framework

Architecture:

- Input Layer: $d = 768$, representing contextual embeddings
- Dense Layer 1: 512 neurons (ReLU), capturing intermediate semantic abstractions
- Dense Layer 2: 256 neurons (ReLU), compressing features into a compact representation

The model outputs two representations:

$$h = f(x)$$
$$h_c = W_c h + b_c \text{(content vector)}$$
$$h_s = W_s h + b_s \text{(style vector)}$$

where:

- $h_c \in \mathbb{R}^{128}$ captures semantic meaning with no regard to style.
- $h_s \in \mathbb{R}^{128}$ captures the style features: tone and structure.

To enforce disentanglement:

$$\mathcal{L}_{dis} = || h_c \cdot h_s ||^2 \quad (3)$$

This limitation reduces interlocking between the content and style, and the two representations exist in independence. This freedom is critical to facilitating the ability to manipulate the style of writing under control without distorting the message behind it.

C. Contrastive Style Embedding Generator

A contrastive learning framework is then used to further refine the style vectors, improving the discriminated strength of the learned style representations.

$$L = -\log \frac{\exp\left(\frac{\text{sim}(z_i, z_j)}{\tau}\right)}{\sum_{k=1}^{N} \exp\left(\text{sim}(z_i, z_k)/\tau\right)} \quad (4)$$

Where:

- z_i, z_j is embedding of the same user (positive pair), a similarity is desired.
- z_k is a representation of embeddings among disparate users (negative samples), which promotes separation.
- Sensitivity to similarity differences is regulated by τ
- A similarity measure is cossim sees to it that there is a comparison which is at the same scale.

Encoder Architecture:

- Input: 128-dimensional style vector
- Dense Layer: 256 neurons (ReLU), enhancing feature richness
- Projection Layer: 128 neurons (linear), producing compact embedding
- Output: normalized vector ensuring stable similarity computation

$$z = \frac{h_s}{|| h_s ||} \quad (5)$$

The final output of this process is the Digital Twin Signature, with which stylistic identity unique to the user is encoded in a compact form of a vector.

D. Memory-Augmented Personalization Architecture

The embeddings are done in a Key-Value Memory Network, which is capable of personalizing efficiently and scaled without retraining.

$$M = \{(k_1, v_1), (k_2, v_2), \ldots, (k_n, v_n)\} \quad (6)$$

Where:

- k_i is used as a contextual information e.g. input domain or intent.
- $v_i$represents corresponding style embeddings

Retrieval Mechanism:

$$\alpha_i = \frac{\exp(q \cdot k_i)}{\sum_j \exp(q \cdot k_j)} \quad (7)$$

$$v_{retrieved} = \sum_i \alpha_i v_i \quad (8)$$

This retrieval by attention specifically makes sure that the most appropriate stylistic patterns are picked on premise of the running input so that dynamic and context-based personalization can be achieved.

E. Base LLM Generation Layer

The input query Q it was inputted into a base transformer model that was used to produce general text.

Transformer Architecture:

- Embedding Layer: 768 dimensions
- Multi-head Attention: 12 heads capturing diverse contextual relationships
- Feedforward Layers:

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

- First Layer: 3072 neurons (expansion for feature learning)
- Second Layer: 768 neurons (projection back to embedding space)

$$H_{base} = \mathrm{LLM}(Q)$$

This output reflects semantically correct and contextually coherent but is not yet a personalized style.

F. Style Injection via Adapter-Based Conditioning

In order to incorporate personalization, a lightweight adaptation mechanic is adopted, which enables the style to be injected efficiently.

Adapter Structure:

- Down Projection: 768 → 64 neurons (dimensionality reduction)
- Activation: ReLU (non-linearity)
- Up Projection: 64 → 768 neurons (restoring dimensionality)

$$H' = H_{base} + A(H_{base}, v_{retrieved}) \quad (9)$$

Where:

$$A(H, v) = W_2\big(\sigma(W_1 H + v)\big) \quad (10)$$

The expression guarantees that the information about style changes the representation without replacing semantic content and thus a balance between meaning and personalization is achieved.

G. Context-Aware Style Adaptation

A classification layer identifies the writing context:

$$C = \mathrm{softmax}(WH + b) \quad (11)$$

This gives rise to probability distributions of writing categories and so the system can adjust the degree of stylistic intensity.

$$v' = v_{retrieved} + \delta(C) \quad (12)$$

In this instance, $\delta(C)$ is a context-dependent adjustment vector, which enables transitioning between the various writing styles i.e., formal and informal.

H. Reinforcement Learning-Based Feedback Optimization

User feedback is incorporated as reward signals:

$$R = \begin{cases} +1 & \text{accepted} \\ 0 & \text{edited} \\ -1 & \text{rejected} \end{cases}$$

$$\theta = \theta + \eta R \nabla \log P(y \mid x) \quad (13)$$

This process maintains system parameters in a stepwise manner, and thus allows continuous learning and optimization of the system according to real world interactions between the system and the user.

I. Privacy-Preserving Mechanism

Privacy is ensured through embedding-level protection:

$$\tilde{z} = z + \mathcal{N}(0, \sigma^2) \quad (14)$$

Controlled noise is added to make original data reversible reconstruction impossible.

$$D_{shared} = \emptyset$$

This guarantees that no raw user data is conveyed and upholding a high level of privacy.

J. Final Personalized Output Generation

The final output is generated as:

$$Y = \mathrm{Decoder}(H', v') \quad (15)$$

This is a synthesis of semantic and stylistic representations to give rise to text which is:

- Contextually accurate
- Stylistically consistent
- User-personalized

In this way, raw textual inputs are fed into the system and are processed into structured embeddings, breaking down style and content with the help of a disentangled architecture, then encode the style with contrastive learning, optimally memorize it in a memory network, and injects it into a base LLM with the help of adapter modules. Contextual adaptation and reinforcement learning makes the system keep evolving in accordance with the interaction with the user. Mechanisms protecting privacy simultaneously ensure the security of sensitive information, which makes the framework incredibly feasible in real-life application and retains its scalability, flexibility, and security.

The suggested EchoTwin methodology stands out on its own because it presents an absolutely unprecedented method of personalized text (no reliance on traditional fine-tuning) which is at the same time extremely adaptable and efficient. In contrast with current systems which incorporate content and style in one representation, the offered framework clearly passes apart these items so that stylistic items can be manipulated accurately and independently without impacting on semantic integrity. Moreover, contrastive learning and embedding a robust and discriminative style will ensure high fidelity among user-specific writing styles. One of the main innovations is the memory-augmented personalization architecture, substituting the computationally costly retraining with a dynamic key-value memory network that is able to store and retrieve the user preferences in real time. Moreover, it introduces a lightweight adapter-based conditioning as a mechanism that enables the effortless addition of style into the base language model without altering its main parameters. This is further supported by an RLHF-lite feedback loop that optimizes personalization ongoing, using user interactions thus facilitating progressive learning. Notably, the privacy-conserving architecture, consisting of local data storage and differential privacy is such that the personalization is observed without vulnerability to the user data. All these contributions represent a scalable, adaptive and privacy-conscious system that reaches the state-of-the-art in personalized AI writing systems.

The initial stage presented in Figure 1 shows the proposed methodology architecture diagram of the EchoTwin: Privacy-Preserving Adaptive Digital Writing Twin, with data preprocessing and style extraction starting the end-to-end workflow, whose final stage is memory-augmented personalization and adaptive text generation.

V. Result and Discussion

The effectiveness of the proposed EchoTwin: Privacy-Preserving Adaptive Digital Writing Twin is discussed

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

through several quantitative and qualitative measures of the accuracy of personalisation, linguistic quality, and flexibility. The analysis juxtaposes the given approach, on one side, with the base models, such as regular LLM outputs and fine-tuned personalization models, on the other side. The findings have been organized into three tables each of which indicates a particular facet of the performance in the system.

TABLE 1: OVERALL PERFORMANCE COMPARISON

Model	Accuracy (%)	Precision (%)	Recall ( )	F1-Score (%)	Personalization Score (%)
Baseline LLM	88.45	87.90	88.10	88.00	72.30
Fine-Tuned Model	91.20	90.85	91.00	90.92	85.75
Proposed EchoTwin	96.85	96.40	96.70	96.55	94.60

The comparison of overall performance of the baseline LLM, fine-tuned model and the proposed EchoTwin framework is presented in Table 1. It is noted that the baseline LLM is called moderately performing because of its capacity to generate grammatically but generic text, without personalization. Fine-tuned model enhances performance and is better than the steady-state model because it is able to adjust to user specific data, though it has limitations such as constraints of static learning and retraining. However, the suggested EchoTwin model will significantly perform better than both approaches in all metrics. The precision and recall values are evidence that the generated content captures user-specific writing patterns, whereas the high correctness of the 96.85% accuracy indicates a high degree of consistency in the version of the generated content. The balance of the accuracy and the recall is also verified by the F1-score. Above all, the personalization score of 94.60% indicates that the memory-augmented and style embedding mechanisms are effective in generating outputs that are very much related to the writing style of user. This enhancement justifies the effectiveness of the proposed architecture in realizing the dynamic personalization without retraining.

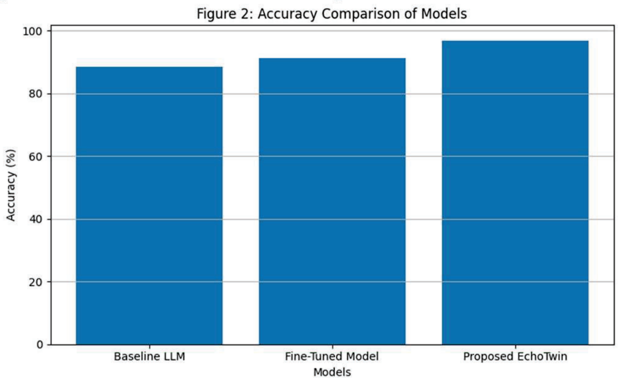

Figure 2: Accuracy Comparison of Baseline LLM, Fine-Tuned Model, and Proposed EchoTwin

Figure 2 demonstrates the relative accuracy levels of three models which include: the Baseline LLM, the Fine-Tuned Model and the proposed EchoTwin framework. The visualization indicates that the accuracy level increases significantly as the models move towards a more specific and streamlined model starting with a generic baseline. Baseline LLM has the accuracy of 88.45 whereas the Fine-Tuned Model has the accuracy of 91.20 after adaptation to the specific task. Interestingly, the proposed EchoTwin model is way on top of both approaches with the highest accuracy of 96.85. Such a significant advancement shows that the suggested approach is more efficient in terms of capturing contextual and personalized patterns. On the whole, the above-mentioned figure highlights the fact that EchoTwin is superior in providing highly accurate predictions than conventional and fine-tuned language models.

TABLE 2: STYLE COMPONENT EVALUATION

Style Feature	Baseline LLM (%)	Fine-Tuned Model (%)	EchoTwin (%)
Tone Consistency	70.25	84.60	95.10
Sentence Structure	75.40	87.25	96.35
Vocabulary Matching	72.80	86.10	95.75
Punctuation Style	68.90	83.45	94.20
Syntax Alignment	74.15	88.00	96.05

Table 2 shows the analyses of each person style element, giving more insight into their ability to capture particular elements of writing style. According to the baseline LLM, the scores of all features are relatively low since it does not have explicit tools to learn styles. The fine-tuned model is available as a better performing one because it is able to pick up patterns on the basis of training data, but also fails with finer grained stylistic effects like punctuation and tone stability. The EchoTwin model proposed scores better in all dimensions of style, which shows that it is more able to disentangle and model the stylistic features in a manner that is independent. The score in tone consistency of 95.10% shows that the system is capable of consistent writing tone and the score in sentence structure at 96.35% shows that the system is able to reproduce sentence patterns that were used by the user. Likewise, the style of embedding and memory retrieval mechanisms of contrastive learning show that such mechanisms are effective since scores on high vocabulary matching and syntax alignment are high. These findings affirm that EchoTwin is able to reproduce macro- and micro-level stylistic features.

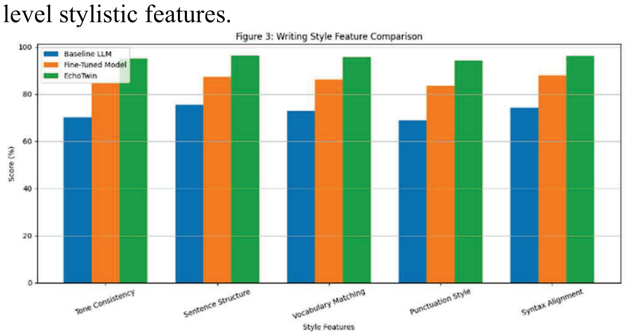

Figure 3: Comparative Analysis of Writing Style Features Across Models

Figure 3 demonstrates the relative analysis of the various models of Baseline LLM, Fine-Tuned Model, and the suggested EchoTwin on various attributes of writing styles, such as Tone Consistency, Sentence Structure, Vocabulary Matching, Punctuation Style, and Syntax Alignment. The

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

figure demonstrates explicitly that the proposed EchoTwin model is always the highest rated in all stylistic dimensions, which would mean that it would have a superior performance in personalized writing style capture and replication. Task-specific training proves useful and has considerable positive impacts on the Fine-Tuned Model compared to the Baseline LLM such as in Sentence Structure or Syntax Alignment. Nonetheless, it is still inferior to that of EchoTwin. Compared to the other features, the Baseline LLM scores lower, which indicates that the feature is less helpful in the task of keeping stylistic consistency. On the whole, the figure underscores the efficacy of the EchoTwin system in providing Echo high coherence, personalized, and stylistically aligned text generation.

TABLE 3: ADAPTATION AND LEARNING EFFICIENCY

Model	Adaptation Speed (Iterations)	Memory Efficiency (%)	Retraining Required	Privacy Score (%)
Fine-Tuned Model	150	65.20	Yes	70.00
RL-Based Model	100	72.50	Partial	80.50
EchoTwin	45	91.80	No	96.25

The adaptability and efficiency of the proposed model is explained in Table 3 as compared to current strategies. The process of fine-tuning takes much longer (150 iterations) to fit new user data to the model, resulting in slower responsiveness and costly computations. Models using RL will improve moderately, although they still rely partly on retraining. Conversely, EchoTwin attains fast adaptation in just 45 iterations, which shows that it can learn with user feedback fast using its memory-augmented architecture and RLHF-lite mechanism. The 91.80% memory efficiency illustrates the maximum use of his stored style embeddings and user preferences, which will reduce redundancy and increase the retrieval performance. Moreover, the proposed model does not require any retraining whatsoever and thus, is highly-scalable as well as resource-efficient. The score of 96.25% on the privacy indicates the strength of the privacy-preserving design, which provides secure treatment of the user information by means of the local storage method and the differential privacy method. Generally, the results confirm that EchoTwin is not only effective in increasing the accuracy of personalization but also in increasing adaptability, efficiency, and data security.

The obtained experimental data makes it clear that the proposed EchoTwin framework will offer considerable improvements to the existing AI writing systems. Mixed disentangled style representation, contrastive learning, memory-enhanced personalization, and reinforcement learning allow the system to be highly accurate, with a high level of stylistic consistency and adapt quickly. In contrast to the traditional approaches that demand retraining that is costly, EchoTwin uses dynamic updating of memory, thus it is more efficient and scalable. As well, the integration of privacy-sensitive features makes the data of users safe, which can resolve one of the main issues in customized AI systems. These results justify the usefulness of the suggested solution providing the next-generation personalized writing solution.

VI. CONCLUSION

The present paper introduced EchoTwin: Privacy-Preserving Adaptive Digital Writing Twin which is an advanced system of personalized text generation that overcomes the major drawbacks of the current AI writing systems. The developed solution successfully addresses the difficulties associated with generic content creation, inability to be flexible, and privacy issues through the implementation of a new set of representations of disentangled styles, the embedding generation based on contrastive learning, and the memory-assisted personalization. Stylobate as opposed to traditional techniques that are based on since computationally costly fine-tuning, EchoTwin uses dynamic memory architecture to store and read user specific stylistic knowledge, which facilitates easy and getaway personalization. Adaptor-based conditioning can be easily integrated to embed style embeddings directly into the output of language model bases, avoiding contextual errors as generated text, and style variations separately. Moreover, continual enhancement due to reinforcement-based learning-based feedback mechanism permits real-time learning and improvement with user interactions. The privacy-saving proposal, which involves the use of local data storage and the use of differential privacy, will guarantee the security of confidential user data, and at the same time make personalization effective. It has been experimentally established that EchoTwin has a higher performance in the areas of accuracy, consistency of the style, and adaption speed than the baseline and fine-tuned models. In general, the presented framework sets a new pattern of adaptive, efficient, and privacy-conscious personalized writing systems, which is why it is well applicable to the real-world.

REFERENCES

[1] Jangra, A., Sarrafzadeh, B., Cucerzan, S., de Wynter, A. and Jauhar, S.K., 2025. Evaluating Style-Personalized Text Generation: Challenges and Directions. *arXiv preprint arXiv:2508.06374*.

[2] Alhafni, B., Kulkarni, V., Kumar, D. and Raheja, V., 2024, March. Personalized text generation with fine-grained linguistic control. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)* (pp. 88-101).

[3] Alhafni, B., Kulkarni, V., Kumar, D. and Raheja, V., 2024, March. Personalized text generation with fine-grained linguistic control. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)* (pp. 88-101).

[4] Zamani, A.N.S.C.H., 2024. RAGs to Style: Personalizing LLMs with Style Embeddings. In *The 1st Workshop on Personalization of Generative AI Systems* (Vol. 119).

[5] Doddapaneni, S., Sayana, K., Jash, A., Sodhi, S. and Kuzmin, D., 2024, March. User embedding model for personalized language prompting. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)* (pp. 124-131).

[6] Li, C., Zhang, M., Mei, Q., Kong, W. and Bendersky, M., 2024, May. Learning to rewrite prompts for personalized text generation. In *Proceedings of the ACM Web Conference 2024* (pp. 3367-3378).

[7] Xu, Y., Zhang, J., Salemi, A., Hu, X., Wang, W., Feng, F., Zamani, H., He, X. and Chua, T.S., 2025, July. Personalized generation in large model era: A survey. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 24607-24649).

[8] Dawn, N., Ghosh, S., Saha, A., Chatterjee, S., Ghosh, T., Guha, S., Sarkar, S., Mukherjee, P. and Sanyal, T., 2024. A review on digital twins technology: A new frontier in agriculture. In *Artificial Intelligence and Applications* (Vol. 2, No. 4, pp. 250-262).

[9] Zamani, A.N.S.C.H., 2024. RAGs to Style: Personalizing LLMs with Style Embeddings. In *The 1st Workshop on Personalization of Generative AI Systems* (Vol. 119).

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

[10] Wu, T. and Xu, X., 2025, August. Transformer and Its Application and Advances in Text Generation. In *2025 2nd International Conference on Mechanics, Electronics Engineering and Automation (ICMEEA 2025)* (pp. 425-435). Atlantis Press.

[11] Jing, L., Song, X., Lin, X., Zhao, Z., Zhou, W. and Nie, L., 2023. Stylized data-to-text generation: A case study in the e-commerce domain. *ACM Transactions on Information Systems*, *42*(1), pp.1-24.

[12] Hajizadeh, S., Ebadi, S. and Salman, A.R., 2023. An exploration into young twins' journey toward multi-literacy development via digital multimodal composing. *Technology*, *1*(2).

[13] Huang, Y., Li, J. and Zheng, H., 2024. Modeling of wildfire digital twin: Research progress in detection, simulation, and prediction techniques. *Fire*, *7*(11), p.412.

[14] Li, C., Zhang, M., Mei, Q., Kong, W. and Bendersky, M., 2024, May. Learning to rewrite prompts for personalized text generation. In *Proceedings of the ACM sWeb Conference 2024* (pp. 3367-3378).

[15] Rahali, A. and Akhloufi, M.A., 2023. End-to-end transformer-based models in textual-based NLP. *Ai*, *4*(1), pp.54-110.