

Physics-Enhanced LSTM Modeling of Electric Double-Layer Transistors

Haoran Yang^{*†}, Shuaidi Zhang^{*†}, Yue Zhao^{*†}, and Ling Li^{*†*}

^{*} State Key Lab of Fabrication Technologies for Integrated Circuits,
Institute of Microelectronics, Chinese Academy of Sciences,
Beijing 100029, China

[†] University of Chinese Academy of Sciences,
Beijing 100049, China

^{*}Corresponding author: lingli@ime.ac.cn

Abstract—This paper proposes a PELSTM framework for behavioral modeling of EDLTs. By embedding physical priors into both the LSTM and MLP modules, the model captures threshold voltage evolution under electric field modulation and accurately describes steady and transient characteristics. Experimental results demonstrate that the proposed framework reproduces the hysteretic I-V characteristics and transient current responses of practical EDLTs with high fidelity. In addition, a PELSTM-based reservoir computing system is constructed for MNIST handwritten digit classification, achieving an average accuracy of 88.4%. The proposed approach offers an effective and physically interpretable solution for modeling complex ion–electron coupled devices and other memristive systems.

Index Terms—electric double-layer transistor, physics-enhanced neural network, long short-term memory, reservoir computing

I. INTRODUCTION

Electric double-layer transistors (EDLTs) have attracted considerable attention due to their high interfacial capacitance, low-voltage operation, and inherent synaptic dynamics, making them promising candidates for neuromorphic computing [1]. However, coupled processes of ion migration, charge redistribution, and interfacial capacitive coupling pose significant challenges for the development of accurate and efficient compact models for EDLTs [2]. To address these challenges, this work adopts a physics-enhanced neural network (PENNN) approach for modeling. Conventional artificial neural networks (ANNs), whose outputs depend solely on current inputs, are unable to capture the intrinsic history dependence and memory effects of the device, thereby limiting their ability to accurately reproduce its dynamic behavior. Recurrent neural networks (RNNs) provide a potential solution to this limitation. Within an RNN, the hidden state (h) primarily serves to retain historical information and can therefore be regarded as a memory module. Building upon this mechanism, the long short-term memory (LSTM) network introduces an additional and independent cell state (c), which enables the network to capture critical information at specific time steps and preserve it over extended periods. Theoretically, LSTM networks are capable of approximating arbitrary nonlinear dynamical systems. Since EDLTs are inherently nonlinear dynamical systems whose hysteresis characteristics depend

strongly on the history of applied gate bias, a long short-term memory (LSTM) network is well-suited for capturing such temporal dependencies. Accordingly, a Physics-Enhanced Long Short-Term Memory (PELSTM) framework is proposed, which integrates physics priors into both the LSTM and multilayer perceptron (MLP) modules.

II. METHODS

A. TCAD Model and Threshold Voltage Drift

In this study, SILVACO TCAD is employed to simulate an ion-conducting electrolyte-gated electric double-layer transistor and to analyze threshold voltage shifts induced by ion migration. The simulated structure [3] is shown in Fig. 1(a). In the simulations, all electrodes are assumed to be Ti. The channel layer is a-IGZO, while the electrolyte layer is TaO_x, where oxygen vacancies serve as the active ionic species. The ion concentration in the electrolyte is 10¹⁸ cm⁻³, the electrolyte thickness is 50 nm, and the channel length is 2 μm. The ion transport process is described using an ion drift-diffusion model. Considering that a-IGZO is an n-type semiconductor, only the corresponding acceptor-like states were taken into account in the simulation.




Fig. 1. (a) Schematic diagram of the simulation structure. (b) Simulated transient response to a V_g pulse. (c) Negative shift of the transfer curves ($V_{ds} = 0.1$ V) over time, with the extracted V_{th} shown in the inset.

To clarify the effect of ion migration in the electrolyte on the threshold voltage (V_{th}) under gate-field modulation, transient and steady simulations are performed using the TCAD model to obtain the current response of the EDLTs to a single voltage pulse of 3 V with a duration of 30 s, as shown in Fig. 1(b). The transfer curves of the device are swept at $t = 1, 5, 10, 15, 20$, and 25 s, and the corresponding threshold

voltages are extracted, as shown in Fig. 1(c). V_{th} exhibits a negative shift as the electric double layer forms at the electrolyte/channel interface.

B. Model Architecture

In this work, two physics-guided design strategies are introduced to enhance the modeling capability of the proposed framework. First, the hidden state (h) output by the LSTM network is employed to characterize the evolution of V_{th} under the influence of the gate electric field. Second, physical prior knowledge is incorporated into the neural network through feature engineering. The architecture of the proposed PELSTM model is illustrated in Fig. 2.




Fig. 2. Architecture of the proposed PELSTM.

The model takes V_{pre} , V , r , and L as inputs, and outputs the drain current I flowing through the EDLTs. Specifically, the hidden state generated at time step $t - 1$ is passed to the LSTM network as an additional input at time step t . This recurrent mechanism enables the model to capture temporal dependencies and characterize the intrinsic memory effects of EDLTs. The PELSTM model consists of two modules [4]. The first is an LSTM module, which is designed to capture the evolution of the internal state associated with the input voltage sequence. Since EDLTs exhibit V_{th} drift during bidirectional scanning of the transfer characteristics, this internal state can be regarded as encoding the evolution of V_{th} . The second is an MLP module, which generates the output current based on the internal state and the remaining input features.

$$h_t = f_{LSTM}(V_t, V_{t-1}, r, h_{t-1}), \quad (1)$$

$$I_t = g_{MLP}\left(V_t - h_t, (V_t - h_t)^2, \frac{1}{L}\right), \quad (2)$$

where f_{LSTM} and g_{MLP} denote the functional mappings implemented by the LSTM and MLP modules, respectively. In (1), V_{t-1} , also denoted by V_{pre} , represents the input voltage at the previous time step and serves as one of the inputs to the LSTM module. V_t represents the input voltage at the current time step, L denotes the channel length of the device, and

TABLE I
HYPERPARAMETER SETTINGS OF THE PELSTM MODEL

Parameters	Quantity
Model configuration	(20, 64, 32, 8)
LSTM parameters	2,000
MLP parameters	5,041
Learning rate	0.001
Training epochs	1,500

r is the gate-voltage sweep rate. In addition, h_{t-1} and h_t represent the hidden states at the previous and current time steps, respectively. Specifically, the LSTM module comprises a single layer with 20 hidden units, whose outputs form the 20-dimensional hidden state h shown in Fig. 2. The MLP module contains three hidden layers with 64, 32, and 8 neurons, respectively. The above architecture yields a total of 2000 parameters for the LSTM module and 5041 parameters for the MLP module. The model is implemented in Python using the open-source deep learning framework PyTorch, and the detailed hyperparameter settings are summarized in Table I.

The mean squared error (MSE) is used as the loss function during training and as an evaluation metric during testing. It is expressed as

$$MSE = \frac{1}{N} \sum_{i=1}^N (I_i - \hat{I}_i)^2, \quad (3)$$

where $\hat{I}_i$ denotes the predicted current at the i th time step, I_i denotes the corresponding ground-truth current, and N is the number of time steps in the sequence.

C. Workflow

Fig 3 illustrates the PELSTM workflow, which comprises several key steps.




Fig. 3. Workflow of PELSTM.

Data Collection. The experimental I-V data of TaO_x/IGZO-based EDLTs with different channel lengths and different gate-voltage sweep rates are collected, as shown in Figs. 4(a) and 4(b), respectively. In addition, the current-response data of an EDLT with $L = 2 \mu\text{m}$ and $r = 0.67 \text{ V/s}$ under voltage-pulse stimulation are used as the training set, as shown in Fig. 4(c).

Feature Engineering. Physical prior knowledge is incorporated into the MLP module based on the current-voltage




Fig. 4. Datasets used for the PELSTM network: I-V characteristics of EDLTs with (a) different channel lengths and (b) different gate-voltage sweep rates. (c) Current-response sequence of an EDLT with $L = 2 \mu\text{m}$ and $r = 0.67 \text{ V/s}$ under voltage-pulse stimulation.

characteristics of thin-film transistors. The drain current in the linear and saturation regimes is expressed as

$$I_{\text{ds}} = \mu C_i \frac{W}{L} \left[(V_{\text{gs}} - V_{\text{th}}) V_{\text{ds}} - \frac{V_{\text{ds}}^2}{2} \right], \quad (4)$$

and

$$I_{\text{ds}} = \mu C_i \frac{W}{2L} (V_{\text{gs}} - V_{\text{th}})^2, \quad (5)$$

respectively. Equation (4) applies to the linear regime, where $V_{\text{ds}} < V_{\text{gs}} - V_{\text{th}}$, whereas (5) applies to the saturation regime, where $V_{\text{ds}} \geq V_{\text{gs}} - V_{\text{th}}$. Here, μ is the field-effect mobility, C_i is the gate capacitance per unit area, W and L are the channel width and channel length, respectively, and V_{gs} and V_{ds} denote the gate-source voltage and drain-source voltage, respectively. Equations (4) and (5) indicate that the drain current depends on the voltage overdrive $V_{\text{gs}} - V_{\text{th}}$, its square, and the reciprocal of the channel length. Accordingly, the LSTM hidden state h is treated as a latent representation of the temporal evolution of V_{th} . The terms $(V - h)$, $(V - h)^2$, and $1/L$ are therefore constructed as physics-informed features and supplied to the MLP module for drain-current prediction.

Model Training. The current values are first transformed using the base-10 logarithm. After the data are divided into training and test sets, min-max normalization parameters are calculated exclusively from the training set and subsequently applied to both sets. The model is trained using the Adam optimizer with a learning rate of 0.001 for 1500 epochs.

Testing. The voltage sequences of the test set are fed into the trained model, and the prediction error is evaluated by comparing the predicted current with the measured current. The voltage step size must remain consistent with that used during training to ensure reliable temporal modeling.

III. EXPERIMENTAL RESULTS

A. PELSTM Validation

The predictive performance of the proposed PELSTM model is first evaluated for the I-V characteristics of EDLTs at a fixed drain bias of $V_{\text{ds}} = 0.1 \text{ V}$. Three independent PELSTM models are trained using the I-V data [5] collected from devices with channel lengths of $L = 2, 5, \text{ and } 10 \mu\text{m}$, respectively. For each dataset, 282 odd-indexed samples are used for training and 282 even-indexed samples are used for testing, with overlap at the two endpoints. For the $L = 5 \mu\text{m}$ dataset, the loss reaches convergence after approximately 1200

epochs. The final training and test losses are 6.02×10^{-6} and 8.21×10^{-6} , respectively, and the total training time is about 1 min, as shown in Fig. 5(a). The R^2 values on both the training and test sets exceed 0.999. The MAPE values on the test sets for $L = 2, 5, \text{ and } 10 \mu\text{m}$ are 1.52%, 2.53%, and 3.16%, respectively. These results indicate that the PELSTM model can accurately capture the device-specific hysteretic I-V behavior, benefiting from its ability to model the evolution of internal states.

To examine the contribution of physics-informed feature engineering, an ablation study is performed by removing the $(V - h)$, $(V - h)^2$, and $1/L$ terms. Figs. 5(b) and 5(c) compare the predictions of the PELSTM models with and without physical priors against the corresponding ground-truth currents on linear and semilogarithmic scales, respectively. The physically informed PELSTM model reduces the test-set MAPE from 6.16% to 2.53%, compared with its counterpart without physical priors, confirming that incorporating relevant physical knowledge improves model accuracy.




Fig. 5. (a) Model predictions for EDLTs with $L = 2, 5, \text{ and } 10 \mu\text{m}$. Comparison of the (b) linear-scale and (c) logarithmic-scale results on the $L = 5 \mu\text{m}$ dataset, with $(V - h)$, $(V - h)^2$, and $(1/L)$ removed from PELSTM.

The proposed PELSTM model is further evaluated for its ability to capture the dependence of the output current on the gate-voltage sweep rate. The I-V data collected from a device with a channel length of $2 \mu\text{m}$ at sweep rates of 0.67, 1.18, and 2.00 V/s are jointly used for model training, whereas the I-V data obtained at a sweep rate of 0.87 V/s are used for testing. V_{ds} is maintained at 0.1 V throughout the measurements. The training results are shown in Fig. 6(a), while the test results are presented in Figs. 6(b) and 6(c). The MAPE values on the training and test sets are 3.06% and 4.41%, respectively. Figs. 6(b) and 6(c) compare the predicted current values with the corresponding ground-truth values on linear and semilogarithmic scales, respectively, demonstrating reasonable prediction accuracy.




Fig. 6. Performance of the PELSTM model trained using data acquired at different gate-voltage sweep rates. (a) Results on the training set. (b, c) Test results at $r = 0.87 \text{ V/s}$ on (b) linear and (c) semilogarithmic scales.

B. PELSTM-Enabled Reservoir Computing

Building on the validated surrogate model described above, PELSTM is further employed to construct a reservoir computing (RC) system [6] for handwritten-digit classification. When PELSTM is treated as a physical emulator, the evolution of its hidden state h provides the nonlinear mapping required for RC. With inherent nonlinearity, a high-dimensional state space, and fading memory, PELSTM serves as a suitable surrogate reservoir for real EDLTs.

Fig. 7(a) shows a 28×28 -pixel image of the handwritten digit “8” and the corresponding preprocessing procedure. Each image is reshaped into a 1×784 vector and linearly mapped to the voltage range of -2.0 – 5.0 V, thereby converting the static image into a temporal gate-voltage sequence. The trained PELSTM network is then used as a virtual device, whose hidden state h provides the nonlinear transformation from the gate-voltage sequence to the intermediate-state sequence, as shown in Fig. 7(b). A modulus operation is subsequently applied to h to obtain a 1×784 reservoir-state vector. This vector is fed into a fully connected layer with 128 neurons, followed by a 10-neuron output layer for digit classification. The model is trained using 20,000 handwritten-digit samples and evaluated on 1,000 test samples. During inference, a winner-take-all (WTA) scheme assigns the class corresponding to the maximum element of the ten-dimensional output vector. The predicted labels are compared with the ground-truth labels to calculate the overall classification accuracy. As shown in Fig. 7(c), the constructed RC system achieves an average classification accuracy of 88.4%, demonstrating its promising capability for temporal information processing and classification.




Fig. 7. (a) Image of the handwritten digit “8” and its preprocessing procedure. (b) Input-voltage pulse sequence generated after preprocessing and the corresponding hidden-state response. (c) Handwritten-digit prediction results obtained using the constructed RC system.

C. Pulse Response Validation

A stage-wise training strategy based on pretrained parameters is adopted. The model is first pretrained for 600 epochs using the I-V data for $L = 2 \mu\text{m}$ and $r = 0.67 \text{ V/s}$, as shown in Fig. 4(a). It is then further trained for 200 epochs using the current-response data shown in Fig. 4(c) for parameter fine-tuning. To evaluate its extrapolation capability, the trained model is tested using a user-defined square-wave voltage-pulse sequence. As shown in Fig. 8, the predicted current largely reproduces the characteristic pulse-response behavior of real EDLTs, in which the current gradually increases with the number of stimuli and exhibits a saturation effect that

is absent from the training set. This result suggests that the model possesses a certain degree of physical interpretability and supports its potential use as a surrogate model for real devices.




Fig. 8. Current response of the PELSTM model under a square-wave voltage input.

IV. CONCLUSION

In this work, a physics-enhanced LSTM (PELSTM) framework is developed for the behavioral modeling of EDLTs. By embedding physical priors into the LSTM and MLP modules, the model learns nonlinear, history-dependent, and hysteretic device dynamics directly from measured data. It accurately reproduces hysteretic I-V characteristics and transient current responses, achieves a training-set MAPE below 4%, and generalizes to operating conditions beyond the training range. The hidden state captures the evolution of V_{th} . Moreover, the PELSTM-based RC system achieves a classification accuracy of 88.4% on MNIST, demonstrating the potential of the framework for modeling EDLTs and other dynamical systems.

REFERENCES

- [1] L. Q. Zhu, C. J. Wan, L. Q. Guo, Y. Shi, and Q. Wan, “Artificial synapse network on inorganic proton conductor for neuromorphic systems,” *Nature communications*, vol. 5, no. 1, p. 3158, 2014.
- [2] J. Bisquert, B. Ilyassov, and N. Tessler, “Switching response in organic electrochemical transistors by ionic diffusion and electronic transport,” *Advanced Science*, vol. 11, no. 36, p. 2404182, 2024.
- [3] R. Fang, X. Li, K. Ren, W. Zhang, H. Xu, L. Wang, and D. Shang, “Improved dynamic characteristics of oxide electrolyte-gated transistor for time-delayed reservoir computing,” *Applied Physics Letters*, vol. 124, no. 5, 2024.
- [4] A. S. Lin, P.-N. Liu, S. Pratik, Z.-K. Yang, T. Rawat, and T.-Y. Tseng, “Rram compact modeling using physics and machine learning hybridization,” *IEEE Transactions on Electron Devices*, vol. 69, no. 4, pp. 1835–1841, 2022.
- [5] R. Fang, S. Wang, W. Zhang, K. Ren, W. Sun, F. Wang, J. Lai, P. Zhang, X. Xu, Q. Luo *et al.*, “Oxide-based electrolyte-gated transistors with stable and tunable relaxation responses for deep time-delayed reservoir computing,” *Advanced Electronic Materials*, vol. 10, no. 4, p. 2300652, 2024.
- [6] X. Liang, J. Tang, Y. Zhong, B. Gao, H. Qian, and H. Wu, “Physical reservoir computing with emerging electronics,” *Nature Electronics*, vol. 7, no. 3, pp. 193–206, 2024.