Regularized Adaptive Weight Noise Injection-Based Evolutionary Training with Generative AI for Industrial Dye Recipe Optimization

Ridha El Hamdi
Laboratory of Advanced Technologies for Medicine and Signals, National Engineering School of Sfax, University of Sfax, Tunisia | Digital Research Center of Sfax, Technopole of Sfax, Tunisia
hamdi.ridha@univgb.tn (corresponding author)

Hana Charaabi
Laboratory of Advanced Technologies for Medicine and Signals, National Engineering School of Sfax, University of Sfax, Tunisia | Digital Research Center of Sfax, Technopole of Sfax, Tunisia
hana.charaabi@enis.tn

Mohamed Njah
Laboratory of Advanced Technologies for Medicine and Signals, National Engineering School of Sfax, University of Sfax, Tunisia | Digital Research Center of Sfax, Technopole of Sfax, Tunisia
mohamed.njah@enis.tn

Mounir Zaag
Industrial Textiles Company (SITEX), Ksar Hellal, Tunisia
zaag.mounir@sitex.com.tn

Received: 14 May 2025 | Revised: 1 July 2025 and 11 July 2025 | Accepted: 13 July 2025

ABSTRACT

This study introduces the Regularized Adaptive Weight Noise Injection-Based Evolutionary (RAWE) training approach, enhanced by generative Artificial Intelligence (AI), to optimize the dye recipe formulation in industrial textile manufacturing. RAWE integrates a self-adaptive evolutionary strategy with adaptive weight noise injection, dynamically balancing the exploration and exploitation during model training. A key innovation of RAWE is its use of generative AI to synthesize high-quality, domain-specific data, addressing the challenge of limited historical dyeing records. This synthetic data generation significantly improves the model generalization and robustness, enabling more accurate and reliable predictions in real-world industrial settings. The effectiveness of RAWE is demonstrated through its deployment in a real-world textile dyeing automation system, where it achieves significant improvements in dye recipe optimization. The results show that RAWE reduces the material waste, minimizes the production costs, and enhances the color consistency compared to traditional methods. By combining generative AI with adaptive evolutionary training, RAWE offers a scalable and practical solution for complex industrial processes, aligning with the latest advancements in automated Machine Learning (ML) and AI-driven optimization.

Keywords-generative AI; generative adversarial networks; variational autoencoders; evolutionary algorithms; weight noise injection; adaptive regularization; industrial dyeing optimization

I. INTRODUCTION

In textile manufacturing, the dye recipe formulation is a critical step, particularly for denim producers, such as SITEX in Tunisia, as it directly affects the shade accuracy, product quality, and operational efficiency. Traditionally, experts rely on trial-and-error to adjust the recipe parameters, which often leads to fabric waste, high energy consumption, and delays when the desired shade is not achieved in early dyeing runs.

ML provides a data-driven solution by leveraging historical production data for recipe prediction. However, applying ML to this domain involves challenges, such as high-dimensional input variables, limited labeled data, and the need for reliable generalization under diverse production conditions [1]. The advances in Evolutionary Algorithms (EAs), noise-based regularization techniques, and generative AI have shown promise in addressing these issues [1, 2].

Injecting noise into model weights helps stabilize the training and reduce the overfitting [3], while Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) generate synthetic data to enrich small datasets [4, 5]. Previous studies have applied Support Vector Machines (SVMs) and Random Forests (RF) for dye recipe prediction [1-6], but these methods often require large labeled datasets. Deep learning models, such as CNNs and RNNs, are capable of modeling the nonlinear dyeing processes [7, 8], but are computationally demanding and vulnerable to overfitting when the data are sparse [7].

Generative AI, including parameterized GANs, has improved the model robustness in fields, such as medical imaging [9], and is increasingly used in textiles for synthetic color generation [10, 11]. Regularization is also critical for generalization. Weight noise injection, which adds stochastic Gaussian noise to weights during training, improves convergence and robustness compared to methods, like weight decay [12, 13], and promotes flatter minima for better model stability [14]. In addition, evolutionary strategies allow a broader exploration of the solution space than gradient-based optimizers like Adam. These strategies have proven effective in optimizing deep networks [15] and in scheduling textile production processes [16].

This paper introduces the RAWE training approach. It is a novel framework that combines adaptive weight noise regularization, VAE-based synthetic data generation, and evolutionary optimization within a unified training loop. Unlike existing hybrid models, RAWE addresses the inverse modeling problem in textile dyeing by predicting input process parameters based on desired output shades. This formulation is rarely explored in the contemporary literature, which typically focuses on forward prediction. The main innovation of RAWE lies in its integration of inverse regression, regularized evolutionary search, and synthetic data augmentation to improve both generalization and robustness.

RAWE was evaluated using real-world production data from SITEX, a textile manufacturer. The results showed significant reductions in shade mismatches, fabric waste, and manual corrections, confirming its effectiveness for industrial dye recipe automation.

II. RAWE TRAINING METHODOLOGY

The RAWE training approach integrates cutting-edge generative AI techniques with adaptive evolutionary strategies to enhance the automation and precision of textile dyeing processes.

The methodology is composed of four key components:

- GAN for data augmentation, enabling the generation of realistic synthetic data to address the scarcity of industrial datasets.
- VAE for latent space modeling and exploration, improving the model's ability to generalize to unseen conditions.
- EA-based ML training, replacing traditional gradient descent with population-based optimization for better convergence in high-dimensional, noisy environments.
- Weight noise injection as a regularization technique, promoting generalization and preventing overfitting by introducing controlled stochasticity into the training process.

The following subsections provide a detailed explanation of the major components of the RAWE methodology, including synthetic data generation, evolutionary training, adaptive optimization, and system deployment.

A. Synthetic Data Generation

Given the limited availability of the labeled datasets in textile dyeing, the presented study employs GANs and VAEs to create synthetic training samples.

1) GANs for Data Augmentation

The generator G learns to produce synthetic fabric properties X_{synth} from random noise z, while the discriminator D classifies whether the input is real or synthetic:

$$X_{synth} = G(z;\theta_G) \quad (1)$$

where θ_G are the learnable parameters of the generator.

The discriminator loss encourages it to distinguish between the real and fake samples:

$$L_D = -\mathbb{E}_{X\sim p_{real}}[Log D(X)] - \mathbb{E}_{z\sim p_z}[Log(1 - D(G(z)))] \quad (2)$$

where $X\sim p_{real}$ represents the draw real samples from the true data distribution, $z\sim p_z$ are the draw latent variables (from a standard normal distribution), $D(X)$ is the discriminator's prediction on real data, and $D(G(z))$ represents the discriminator's prediction on fake data generated by $G(z)$

The generator is trained to maximize the discriminator's belief that the generated data are real, hence, minimizing the aforementioned loss:

$$L_G = -\mathbb{E}_{z\sim p_z}[Log D(G(z))] \quad (3)$$

This adversarial training ensures that the generated synthetic samples match the distribution of the real fabric properties.

2) VAE for Latent Space Exploration

A VAE learns a compressed, probabilistic latent representation z of input data X via an encoder-decoder structure. The encoder approximates the true posterior with:

$$q(z|X) = \mathcal{N}(\mu(X), \sigma^2(X)) \quad (4)$$

The decoder reconstructs X from z, yielding $\hat{X}$. The VAE loss is a combination of:

Reconstruction error: how well $\hat{X} \approx X$. KL-divergence: regularizes the latent space to follow the prior $p(z)=N(0, I)$.

$$L_{VAE} = \mathbb{E}_{q(z|X)}[\|X - \hat{X}\|^2] + D_{KL}(q(z|X)||p(z)) \quad (5)$$

This enables the controlled sampling of realistic input conditions and supports generalization.

B. Evolutionary Training with Weight Noise Injection

The RAWE training employs an evolutionary strategy, where the mutation-based optimization adapts the weight perturbation dynamically.

1) Adaptive Noise Injection for Regularization

Instead of relying on static weight decay, adaptive Gaussian noise is injected into the model weights to promote generalization and prevent overfitting. Specifically, noise η is added element-wise to the weights W, where each component of η is sampled from a normal distribution with mean zero and variance σ^2:

$$W' = W + \eta\ ; \eta \sim \mathcal{N}(0, \sigma^2) \quad (6)$$

where σ is a scalar that controls the magnitude of the injected noise and is adaptively updated during training based on the gradient of the fitness function:

$$\sigma_{t+1} = \sigma_t \cdot exp(\frac{\partial f}{\partial \sigma}) \quad (7)$$

This approach ensures a broad exploration of the solution space during early training and fine-tuned adjustments in later stages, making the learning process more robust in high-dimensional, noisy environments.

2) Evolutionary Algorithm for Model Selection

Instead of using gradient descent, the model parameters were optimized through mutation and selection:

- Mutation: Generation of n variations of the model parameters Wi with noise perturbation.
- Evaluation: Computation of a fitness score based on validation loss.
- Selection: Retention of the top k models for the next iteration.

The fitness function is defined as:

$$f(W) = -MSE(Y, \hat{Y}) \quad (8)$$

where Y is the true dyeing output and $\hat{Y}$ is the predicted dyeing outcome.

To guide model learning, RAWE minimizes a composite loss that balances the dyeing prediction accuracy and latent space regularization. The final training loss function is defined as:

$$L_{RAWE} = \underbrace{MSE(Y, \hat{Y})}_{Recipe\ Accruacy} + \underbrace{\lambda D_{KL}(q(z|X)||p(z))}_{Laten\ Regularization} \quad (9)$$

$MSE(Y, \hat{Y})$ penalizes the prediction errors. $D_{KL}(q(z|X)||p(z))$ encourages the latent distribution to follow a standard prior. λ is a weighting factor that balances the trade-off between the accurate prediction and latent space regularity.

C. Adaptive Optimization for Dyeing Formulation

The RAWE approach learns an inverse mapping from the textile properties to the optimal input conditions. Instead of predicting the outcomes, the model generates the ideal fabric processing parameters for a desired color output Y^*:

$$X^* = arg\ \min_X \|f(X) - Y^*\| \quad (10)$$

Using Covariance Matrix Adaptation Evolution Strategy, (CMA-ES) (gradient-free optimization), the model iteratively adjusts the fabric treatment parameters to meet the target dyeing characteristics.

III. EXPERIMENTAL RESULTS AND EVALUATION

The proposed RAWE-based generative AI system was developed and deployed along with SITEX, a major Tunisian industrial textile company specializing in denim fabric production. This collaboration enabled both a rigorous experimentation and real-world validation of the intelligent dyeing formulation system.

A. Industrial Context and Motivation

At SITEX, Indigo rope dyeing is used to impart the desired blue shade to yarns before weaving. Unlike conventional fabric dyeing, this process operates on yarns grouped in cables, making it highly sensitive to input variations, such as the cotton quality, room humidity, and customer-specific washing recipes.

The dyeing operation has long relied on manual expert adjustments to find the optimal recipe for each customer-requested shade, defined by three colorimetric parameters in the CIELAB space, which is a perceptually uniform color space, developed by the International Commission on Illumination, designed to approximate human vision: L (lightness), ranging from 0 (dark) to 100 (light); a, ranging from green (negative) to red (positive); and b, ranging from blue (negative) to yellow (positive).

CIELAB helps precisely define a shade that a customer requests and enables a quantitative comparison between: the requested shade from the customer and actual shade from the dyed fabric. Often, the difference between two colors is measured using:

$$\Delta E = \sqrt{(L_1 - L_2)^2 + (a_1 - a_2)^2 + (b_1 - b_2)^2} \quad (11)$$

where ΔE quantifies how far two shades are from each other in the CIELAB space.

Failures in early dyeing trials often resulted in material wastage, customer rejection, or costly delays.

B. Dataset Structure and Variables

The RAWE model was trained on historical production records from SITEX. The dataset covers over ten years of Indigo rope dyeing operations and includes 338 samples, each corresponding to a real dyeing process. Each sample contains 34 input parameters and 32 output parameters, categorized to reflect the operational and business priorities, as shown in Table I. The inputs include customer shade requirements, cotton characteristics, and process conditions, while the outputs represent the dye recipe parameters.

TABLE I. CATEGORIZED INPUT AND OUTPUT VARIABLES USED IN THE RAWE MODEL

Type	Group	Parameters	No. of parameters
Inputs (total 34)	Customer requirements	L_fabric, a_fabric, b_fabric	3
	Cotton characteristics	Micronaire, El_fabric, Trh_grad, Maturity, Fiber length, unif_idx, S-F_percent, Tenacity, De_sci, De_rd, De_b	11
	Washing recipe	t°_des, tim_des, cn_des, cn_disp, t°_pumice, time_pum, cn_pumice, t°_cellu, tim_cellu, cn_cellu, t°_detr, tim_detr, cn_detr, t°_white, tim_white, cn_white, t°_netr, tim_netr, cn_netr, cn_adouc	20
Outputs (total 32)	Cable shade	L_cable, a_cable, b_cable	3
	Sulfur bath (B1)	Temp_B1, TB_DrC, TB_InitC	3
	SH parameters	SH_DrC, SH_InitC	2
	BC parameters	BC_DrC, BC_initC	2
	GD parameters	GD_DrC, GD_InitC	2
	DSA parameters	DSA_DrC, DSA_InitC	2
	YA parameters	YA_DrC, YA_InitC	2
	Indigo bath	pH_IB, Temp_IB, I_Fr, H_Fr, SH_Fr	5
	Consumption indicators	I_Cons, Init_I_Cons, H_Cons, Init_H_Cons, SH_Cons, Pr_Cons, C_Cons, TB_Cons, B_Cons, K_Cons, K_InitC	11

The data were split into 236 training samples (70%), 34 validation samples (10%), and 68 test samples (20%). To improve generalization, synthetic samples were generated via a VAE and mixed with real data during training in an 80/20 real-to-synthetic ratio. The validation and testing used only real samples to ensure the objective performance evaluation.

C. Training Procedure

All RAWE training experiments were conducted using Kaggle Notebooks (Python 3, up to 30 GB RAM). The RAWE model, with and without generative AI augmentation, was trained on the historical SITEX dataset using an EA for 2000 generations (population size 50, mutation rate 0.1, initial Gaussian noise $\sigma = 0.05$ adaptively tuned).

Synthetic data (20% of the training set) were generated via a VAE with a 5-dimension latent space. Each neural network in the population had two hidden layers (128 and 64 neurons, ReLU activation) and a linear output. The models were trained using MSE loss.

D. Experimental Results and Discussion

The training dynamics of the RAWE algorithm under both conditions (with and without generative AI) were closely monitored. Figures 1 and 2 illustrate the progression of the average Mean Squared Error (MSE) over the first 1300 training generations for RAWE without and with generative AI, respectively. Figures 3 and 4 show the corresponding best MSE values over the same generations.

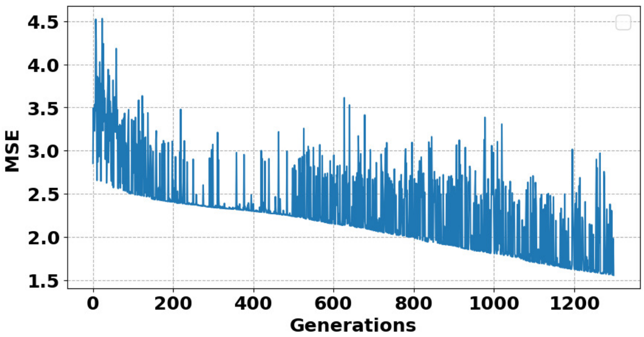

Fig. 1. Average fitness over generations for RAWE without generative AI.

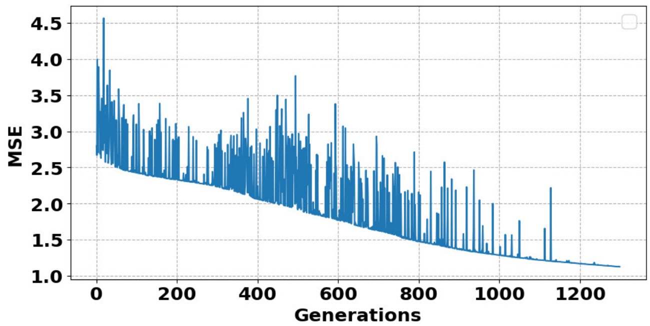

Fig. 2. Average fitness over generations for RAWE with generative AI.

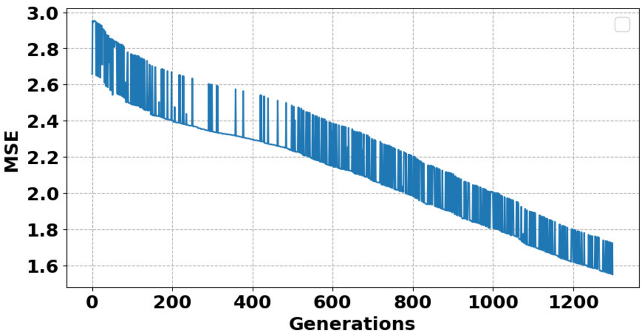

Fig. 3. Best fitness over generations for RAWE without generative AI.

It is observed that the RAWE with generative AI consistently demonstrates a faster and smoother decrease in both the average and best MSE compared to the RAWE without generative AI. The final MSE achieved by the augmented RAWE is significantly lower, indicating a substantial improvement in the model's ability to accurately predict the optimal dye recipes.

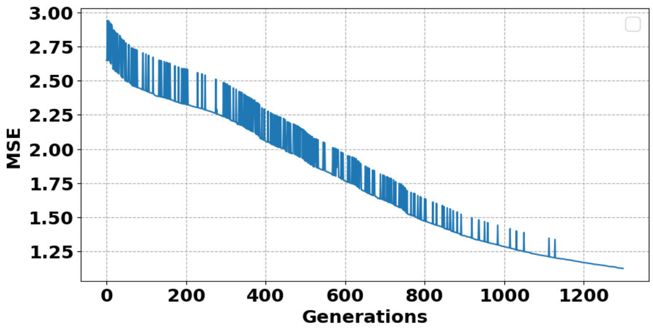

Fig. 4. Best fitness over generations for RAWE with generative AI.

This enhanced accuracy directly supports the central claim of the study: that RAWE, particularly when augmented with generative AI, leads to a reduction in the shade mismatches, material waste, and the time required for manual corrections, thereby offering a more efficient and scalable solution for industrial dye recipe automation. It is important to note that while the models were trained for 2000 generations, only the first 1300 generations are shown in the current study to provide a clearer visualization of the performance progression.

The RAWE model is trained to minimize the loss function (10). However, the primary objective in this industrial application is to match the customer's desired fabric shade (L_fabric, a_fabric, b_fabric). Figures 5-7 illustrate the true and predicted values of L_cable, a_cable, and b_cable on a subset of data, evaluating the model's practical shade prediction performance with generative AI.

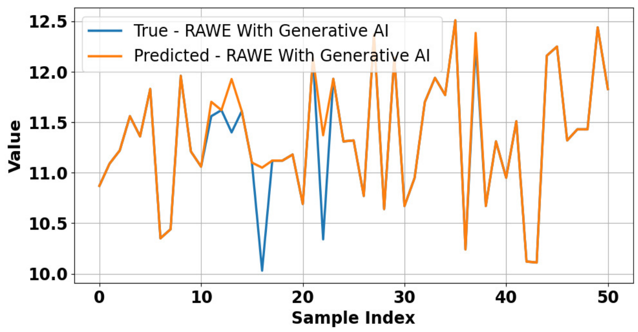

Fig. 5. True versus predicted L_cable values using RAWE with generative AI.

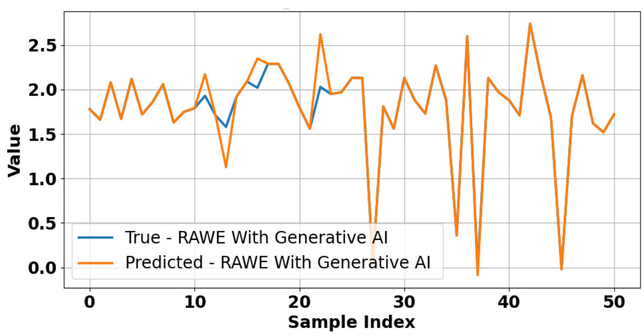

Fig. 6. True versus predicted a_cable values using RAWE with generative AI.

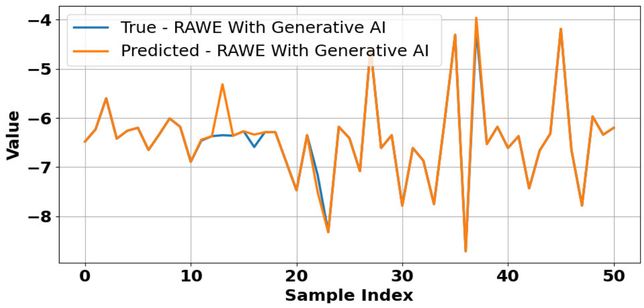

Fig. 7. True versus predicted b_cable values using RAWE with generative AI.

The model generally captures the overall trend of the fabric shade values, with predictions often closely following the true values. However, there are instances where the model slightly overestimates or underestimates the lightness. The close alignment in many regions suggests that RAWE with generative AI can effectively learn the relationship between the input parameters and the resulting fabric shade values.

IV. COMPARISON WITH BASELINE MODELS AND ABLATION STUDY

To further validate RAWE's performance, it was compared with three commonly used regression models: Traditional Artificial Neural Network (ANN) trained using Adam optimizer, Support Vector Regression (SVR), and Random Forest Regressor (RF). All baseline models were trained and tested exclusively on the same real dataset. In contrast, RAWE was evaluated in multiple configurations to assess the contributions of its core components:

- RAWE (Real only): trained solely on real data, without any synthetic samples.
- RAWE (Full model): trained on both real and synthetic data generated via a VAE, with full use of evolutionary training and adaptive noise injection.
- RAWE (No VAE): trained with real data and adaptive training, but without synthetic augmentation.
- RAWE (No Noise): trained with real and synthetic data, but without Gaussian noise injection into model weights.
- RAWE (No Evolution): trained with real and synthetic data and noise regularization, but using standard Adam optimizer instead of evolutionary selection.

All models were evaluated using MSE and Mean Absolute Error (MAE) on the same test set. The primary results are summarized in Table II, while the ablation analysis is detailed in Table III.

TABLE II. EVALUATION OF MODELS

Model	MSE (real scale)	Model (real scale)
RAWE (with generative AI)	61.840	3.93
RAWE (without generative AI)	78.650	4.75
ANN + Adam	95.120	5.45
SVR	108.34	5.77
RF regressor	87.430	5.18

Table II presents the performance of RAWE compared to conventional regression models. RAWE with generative AI achieves the best results, with a final MSE of 61.84 and a MAE of 3.93, significantly outperforming traditional approaches, such as ANN trained with the Adam optimizer, SVR, and RF.

RAWE trained without synthetic data still outperforms all baseline models. This demonstrates that the combination of adaptive weight noise injection and evolutionary optimization alone provides strong generalization capabilities for the inverse modeling task. When generative augmentation is included, the performance improves further, validating the value of synthetic data in addressing limited industrial datasets.

TABLE III. ABLATION STUDY – CONTRIBUTION OF RAWE COMPONENTS

Model variant	Synthetic data	Noise injection	Evolutionary training	MSE	MAE
RAWE (Full Model)	Yes	Yes	Yes	61.84	3.93
RAWE (No VAE)	No	Yes	Yes	77.25	4.62
RAWE (No Noise)	Yes	No	Yes	81.4	4.89
RAWE (No Evolution)	Yes	Yes	No	78.73	4.69

Table III provides an ablation study evaluating the contribution of each RAWE component. The removal of any module (synthetic data, noise regularization, or evolutionary training) leads to degraded performance. RAWE (Full) confirms the synergy between all three components, while RAWE (Real Only) shows that the training strategy itself is effective, even without augmentation. Each component independently improves the predictive accuracy, and their combination achieves optimal performance.

The MSE values during curve training were computed on normalized outputs (range (0, 1)), while the final MSE and MAE scores in Tables II and III are computed after inverse scaling. Thus, they reflect real-world physical units relevant to the dyeing parameters, making them directly interpretable for industrial deployment.

Overall, these results highlight the robustness, modularity, and effectiveness of RAWE for high-dimensional inverse regression tasks, such as automated dye recipe formulation in textile production.

V. CONCLUSION

This paper presented the Regularized Adaptive Weight Noise Injection-Based Evolutionary (RAWE) training approach, augmented with generative Artificial Intelligence (AI), to optimize the dye recipe formulation in the industrial textile domain. RAWE combines a self-adaptive Evolutionary Algorithm (EA) with adaptive Gaussian noise injection, enabling robust training even in high-dimensional, data-scarce scenarios. This synergy enhances both the exploration and convergence during optimization.

The successful deployment of RAWE at SITEX, a leading Tunisian denim manufacturer, demonstrates its practical industrial relevance—significantly reducing the trial-and-error iterations, lowering material waste, and improving the responsiveness to client-specific shade requirements.

While synthetic data generation improves the generalization and mitigates data scarcity, it also introduces potential risks. The over-reliance on synthetic samples may lead to distributional shifts or learning biases, especially in underrepresented or edge-case conditions. To address this issue, future works should explore enhanced validation mechanisms for synthetic inputs, investigate diffusion-based generative models, test RAWE across broader industrial datasets, and use cases to evaluate its scalability and generalizability beyond textile dyeing.

REFERENCES

[1] Md. N. Pervez, W. S. Yeo, L. Lin, X. Xiong, V. Naddeo, and Y. Cai, "Optimization and prediction of the cotton fabric dyeing process using Taguchi design-integrated machine learning approach," *Scientific Reports*, vol. 13, no. 1, July 2023, Art. no. 12363, https://doi.org/10.1038/s41598-023-39528-1.

[2] M. Njah and R. El Hamdi, "A Constrained Multi-Objective Learning Algorithm for Feed-Forward Neural Network Classifiers," *Engineering, Technology & Applied Science Research*, vol. 7, no. 3, pp. 1685–1693, June 2017, https://doi.org/10.48084/etasr.968.

[3] R. M. Zur, Y. Jiang, L. L. Pesce, and K. Drukker, "Noise injection for training artificial neural networks: A comparison with weight decay and early stopping," *Medical Physics*, vol. 36, no. 10, pp. 4810–4818, Oct. 2009, https://doi.org/10.1118/1.3213517.

[4] M. Frid-Adar, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "Synthetic data augmentation using GAN for improved liver lesion classification," in *2018 IEEE 15th International Symposium on Biomedical Imaging*, Washington, DC, Apr. 2018, pp. 289–293, https://doi.org/10.1109/ISBI.2018.8363576.

[5] K. Dorji, S. Jittanon, P. Thanarak, P. Mensin, and C. Termritthikun, "Electricity Load Forecasting using Hybrid Datasets with Linear Interpolation and Synthetic Data," *Engineering, Technology & Applied Science Research*, vol. 14, no. 6, pp. 17931–17938, Dec. 2024, https://doi.org/10.48084/etasr.8577.

[6] N. Tarafder, "Applications of Machine Learning in Textile Industry," *Journal of IoT and Machine Learning*, vol. 1, no. 2, pp. 24–27, Aug. 2024, https://doi.org/10.48001/joitml.2024.1224-27.

[7] N. Ingle and W. J. Jasper, "A review of deep learning and artificial intelligence in dyeing, printing and finishing," *Textile Research Journal*, vol. 95, no. 5–6, pp. 625–657, Mar. 2025, https://doi.org/10.1177/00405175241268619.

[8] M. An, S. Wang, L. Zheng, and X. Liu, "Fabric defect detection using deep learning: An Improved Faster R-approach," in *2020 International Conference on Computer Vision, Image and Deep Learning*, Chongqing, China, July 2020, pp. 319–324, https://doi.org/10.1109/CVIDL51233.2020.00-78.

[9] X. Xiong *et al.*, "A Parameterized Generative Adversarial Network Using Cyclic Projection for Explainable Medical Image Classifications," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing*, Seoul, Korea, Republic of, Apr. 2024, pp. 7310–7314, https://doi.org/10.1109/ICASSP48485.2024.10448260.

[10] L. Schaerf, A. Alfarano, and E. Postma, "ColorwAI: Generative Colorways of Textiles Through GAN and Diffusion Disentanglement," in *Computer Vision – ECCV 2024 Workshops*, vol. 15628, A. Del Bue, C. Canton, J. Pont-Tuset, and T. Tommasi, Eds. Cham: Springer Nature Switzerland, 2025, pp. 137–160.

[11] A. Imtiaz, N. Pathirana, S. Saheel, K. Karunanayaka, and C. Trenado, "A Review on the Influence of Deep Learning and Generative AI in the Fashion Industry," *Journal of Future Artificial Intelligence and Technologies*, vol. 1, no. 3, pp. 201–216, Oct. 2024, https://doi.org/10.62411/faith.3048-3719-29.

[12] X. Wang, L. Kong, and L. Wang, "Estimation of Error Variance in Regularized Regression Models via Adaptive Lasso," *Mathematics*, vol.

10, no. 11, June 2022, Art. no. 1937, https://doi.org/10.3390/math10111937.

[13] P. Züge, C. Klos, and R.-M. Memmesheimer, "Weight versus Node Perturbation Learning in Temporally Extended Tasks: Weight Perturbation Often Performs Similarly or Better," *Physical Review X*, vol. 13, no. 2, Apr. 2023, Art. no. 021006, https://doi.org/10.1103/PhysRevX.13.021006.

[14] A. Orvieto, A. Raj, H. Kersting, and F. Bach, "Explicit Regularization in Overparametrized Models via Noise Injection." arXiv, 2022, https://doi.org/10.48550/ARXIV.2206.04613.

[15] B. Wang, Y. Sun, B. Xue, and M. Zhang, "Evolving deep neural networks by multi-objective particle swarm optimization for image classification," in *Proceedings of the Genetic and Evolutionary Computation Conference*, Prague Czech Republic, July 2019, pp. 490–498, https://doi.org/10.1145/3321707.3321735.

[16] E. Duran, C. Ozturk, and M. A. Ornek, "Combinatorial optimization methods for yarn dyeing planning," *Flexible Services and Manufacturing Journal*, vol. 37, no. 1, pp. 282–319, Mar. 2025, https://doi.org/10.1007/s10696-024-09541-1.