DRCoD: Toward Robust Continual Learning of Diffusion Models for Tire Manufacturing Prototyping

Supplementary Material

9. Traditional Product Development With Physical Prototyping

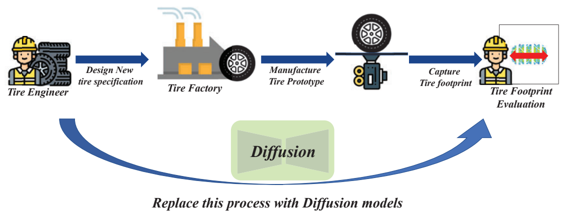

Figure 8. Visualization of the process of new product development in tire industry

In most manufacturing industries, the development of new products involves a prototyping process to validate ideas and identify potential issues in functionality and performance before mass production. This development process is typically iteratively carried out, where physical prototypes are produced 3D printing or injection molding, followed by performance testing and modification of specification for the final product. For example, as illustrated in Fig. 8, in the tire manufacturing industry, engineers first design the tire specifications and then produce prototypes, which are tested on the ground while cameras capture the contact patches. These footprint images are used not only as critical indicators for evaluating tire performance but also to verify whether the designed specifications are visually realized as intended. Specifically, the contact patch area and shape are closely related to load transfer and grip, while the pressure distribution reflects handling stability and localized stress. Based on the footprint results, engineers modify the specifications and, when necessary, produce new prototypes for subsequent validation. However, such physical prototype-based product development consumes labor and economic resources, while also imposing environmental burdens through manufacturing and disposal processes. With recent advances in generative models, particularly diffusion models, it has become possible to generate high-quality images that align with given conditions. Nevertheless, direct application of generative models to product development in the manufacturing industry is not straightforward. In real manufacturing environments, data are provided in an online setting rather than being accessible all at once, as new categories, such as new materials, patterns, and operating conditions, are continuously introduced. To remain effective, models must therefore be capable of continual learning, but this inevitably causes catastrophic forgetting in diffusion models, severely limiting their applicability in practice. To overcome this challenge, this study proposes a novel approach that enables diffusion models to adapt continually while preserving previously acquired knowledge.

10. Comparison with existing continual learning methods

Tab. 6 compares our proposed method with existing continual learning approaches. Generative distillation and C-LoRA are designed considering the structural characteristics of diffusion models. However, these methods neither account for dataset imbalance and data complexity, nor are they typically validated in settings like industrial prototyping, where each condition corresponds to a unique ground-truth image. This requirement differs fundamentally from general text-to-image generation, where no unique ground-truth image exists, as a single prompt can naturally correspond to multiple equally valid outputs (e.g., the prompt "a dog in the grass" may yield images with different poses, colors, or backgrounds). Moreover, C-LoRA requires task-specific parameter storage for each task, and Generative distillation necessitates copying and storing previous models for training. Among conventional approaches, replay-based methods can adjust sampling strategies to address dataset imbalance when storing data in the buffer. However, data-level adjustment does not work for diffusion. To the best of our knowledge, DRCoD is the first continual learning method specifically designed for diffusion models that addresses dataset imbalance.

11. Implementation Details

11.1. Datasets

The dataset was collected from a global tire manufacturing company. For each sample, real physical tire prototypes were produced according to specifications provided in tabular form, and the corresponding footprint images were captured. Accordingly, each data instance consists of a table of tire specifications paired with a footprint image obtained from the corresponding physical prototype. The specification data include not only various design-related variables such as material and thickness, but also test conditions, in-

Category	Method	Designed for Diffusion	No Parameter Storage	Adaptability to complex data
Regularization-based	EWC	✗	✓	✗
	MAS	✗	✓	✗
	Generative Distillation	✓	✗	✗
Replay-based	Experience Replay	✗	✓	✓
	Generative Replay	✗	✓	✓
Parameter Isolation	C-LoRA	✓	✗	✗
Hybrid (Ours)	**DRCoD**	✓	✓	✓

Table 6. Comparison of Methods for Continual Learning

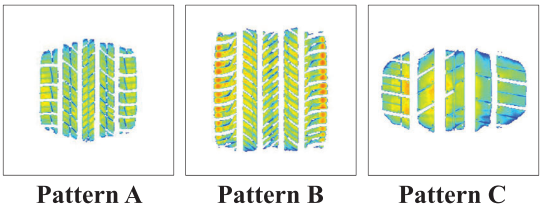

Figure 9. Example footprint images obtained from three tread patterns (Pattern A, Pattern B, and Pattern C)

cluding Test Load and Inflation Air. Importantly, the dataset does not contain any column that directly describes the contact area or shape of the footprint. These variables are interdependent, and more specifically, categorical variables such as pattern determine the type of material, thickness, and structural arrangement, which in turn lead to variations in the shape, area, and pressure distribution of the footprint images. As illustrated in Fig. 9, different patterns result in characteristic footprint shapes. Therefore, in order to generate high-fidelity images, the model must train with newly developed patterns with their corresponding footprint shape and pressure distribution. Since such patterns are continuously introduced in real-world manufacturing, continual learning becomes indispensable for maintaining the model's ability to adapt to new designs without forgetting previously learned ones.

11.2. Loss Weight Scheduler Details

To ensure a gradual and stable transition between tasks, we employ a loss weight scheduler $\psi(\cdot)$ that modulates the influence of regularization throughout training. This strategy is inspired by the warm-up loss technique used in VAE, where the training initially focuses on reconstruction loss to encourage maintaining structural consistency, while the KL divergence loss weight is gradually increased for the latent space structure. Similarly, our method initially emphasizes learning the current task while gradually increasing the weight of regularization to mitigate catastrophic forgetting. The loss weight scheduling follows a cosine-based formulation:

$$\psi(e) = \left(0.5 \times \left(1 - \cos\left(\pi \times \min\left(1.0, \frac{e}{E}\right)\right)\right)\right)^2 \quad (13)$$

where e represents the current epoch and E denotes the total number of epochs. This scheduling method ensures a smooth progression of regularization, preventing interference with the learning of the current task due to constraints. Our results demonstrate that this approach enhances model stability by dynamically adjusting the impact of past tasks, maintaining a balance between retaining past knowledge and adapting to new tasks.

12. Experimental Details

12.1. Continual Learning Evaluation Metrics

We measure the continual learning performance of the model using three key metrics: Average Incremental Quality (AIQ), Average Final Quality (AFQ), and Forgetting Rate (FR).

Average Incremental Quality (AIQ) AIQ evaluates the historical learning performance across tasks. It is computed as:

$$AIQ = \frac{1}{T}\sum_{t=1}^{T}\left\{\frac{1}{t}\sum_{i=1}^{t} m(f^{(t)}, \mathcal{T}^{(i)})\right\} \quad (14)$$

where $m(f^{(t)}, \mathcal{T}^{(i)})$ represents the model's performance on task i after training on task t which represents average quality.

Average Final Quality (AFQ) AFQ quantifies the final performance of the model after completing all learning tasks, which is important in downstream applications. It is computed as:

$$AFQ = \frac{1}{T}\sum_{t=1}^{T} m(f^{(T)}, \mathcal{T}^{(t)}) \quad (15)$$

which measures the model's overall performance across all the tasks.

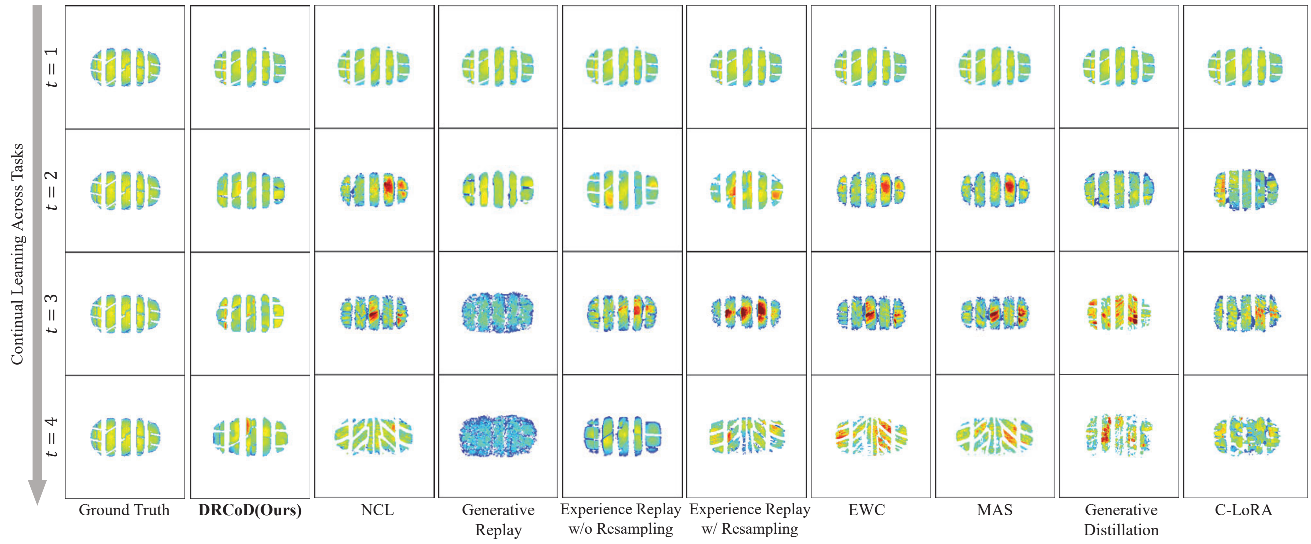

Figure 10. Qualitative comparison of generated tire footprint images of head condition in a continual learning setup

Forgetting Rate (FR) FR assesses how much the model forgets previously learned knowledge when adapting to new tasks. It is defined as:

$$FR^{(t)} = \frac{1}{T-t}\sum_{i=t+1}^{T}\left[m(f^{(i)},\mathcal{T}^{(t)}) - m(f^{(t)},\mathcal{T}^{(t)})\right] \tag{16}$$

where a lower FR indicates that the model effectively preserves past knowledge.

12.2. Qualitative Results of Head Condition

Fig. 10 shows the tire footprint images generated based on the conditions belonging to the head class. Compared with the Ground Truth (GT), DRCoD (Ours) successfully generates footprint images corresponding to the given conditions, not only for the tail but also for the head conditions. Specifically, even after learning Tasks 2-4, it retains knowledge from Task 1 without forgetting. In contrast, NCL, EWC, and MAS lose most of the information, ultimately generating footprint images that correspond to different conditions. Generative Replay and c-LoRA exhibit performance degradation, while Generative Distillation retains information well up to Task 2 but deteriorates as the number of tasks increases. ER w/o resampling, unlike in the tail condition, mitigates forgetting in the head condition. However, it still exhibits unstable generation performance, such as failing to capture the pressure distribution, which is represented by color distribution, after learning Task 3 and 4.