arXiv:2409.01348v1 [cs.CV] 2 Sep 2024

PatternPaint: Generating Layout Patterns Using Generative AI and Inpainting Techniques

Guanglei Zhou*, Bhargav Korrapati†, Gaurav Rajavendra Reddy†, Jiang Hu‡, and Yiran Chen* Dipto G. Thakurta†
*Dept. of Electrical & Computer Engineering, Duke University, Durham, USA
†Intel Corp., Hillsboro, USA
‡Dept. of Electrical & Computer Engineering, TAMU, College Station, USA

***Abstract*—Generation of VLSI layout patterns is essential for a wide range of Design For Manufacturability (DFM) studies. In this study, we investigate the potential of generative machine learning models for creating design rule legal metal layout patterns. Our results demonstrate that the proposed model can generate legal patterns in complex design rule settings and achieves a high diversity score. The designed system, with its flexible settings, supports both pattern generation with localized changes, and design rule violation correction. Our methodology is validated on Intel 18A Process Design Kit (PDK) and can produce a wide range of DRC-compliant pattern libraries with only 20 starter patterns.**

I. INTRODUCTION

Lithography stands as a cornerstone in the advancement of semiconductor technology, presenting significant challenges in the evolution of cutting-edge technology nodes. At the heart of this progression lies the development of Optical Proximity Correction (OPC) [1], [2], [3], [4], [5] elements, hotspot detection[6], [7], [8], [9], [10] and other lithographic strategies, initiated in the early stages of technology development. During this phase, foundational Design Rules (DRs) are established, yet actual layout data remains scarce. Consequently, lithography developers typically resort to traditional test designs or extrapolate from previous technology nodes, resulting in methodologies that are attuned to a narrow spectrum of patterns, neglecting the wider diversity encountered in later stages. As technology transitions into production and more complex designs emerge, numerous unanticipated patterns surface, and often cause systematic defects.

As technology keeps scaling down, the landscape of chip manufacturing has witnessed a dramatic increase [11] in the complexity and volume of design rules required to ensure manufacturability at advanced process geometries. This surge in design rules is largely driven by the need to manage increasing variability, edge placement errors, and interdependencies among rules which can have cascading effects on the manufacturability and reliability of chips. This complexity not only makes manual tracking of rules impractical at below 28nm but also extends challenges to synthetic pattern generations.

Before the rise of machine learning, a number of rule/heuristic-based methods[12], [13], [14] were proposed to generate synthetic layout patterns. However, these heuristic methods often require design rules to be converted into algorithmic constraints, demanding substantial engineering effort during development and adaptation to new technology nodes.

More recently, ML-based methods, leveraging generative models such as GANs, Transformers, TCAEs, and Diffusion models [15], [16], [17], [18] have been proposed with the promise of reduced engineering effort and high pattern diversity. At the time, directly inputting the layout patterns as pixel-based images and regenerating them is deemed almost impossible due to the presence of noise in the generated images. Therefore, most methods [15], [18], [16] take one step of simplification to decompose the layout patterns as a squish[19], [20] representation which consists of a topology matrix and geometry vectors that record the physical information of each row and column. The ML model then only focuses on generating pattern topologies (a blue-print of a layout pattern that only consists of the shape of patterns). With the help of “legalization” (a nonlinear solver) to convert topology into DR clean layout patterns, these methods are successfully used in 1D unidirectional layout pattern generations and 2D relaxed DR settings (min-width, min-spacing, min-area, max-area) layout pattern generation. However, for the industrial standard DR settings where it often involves not just upper bound of width and length but also only allows a set of discrete values of certain rules, the “legalization” then becomes a mixed integer linear programming problem. Its runtime will grow exponentially with the increase of topology size and the fast-growing # of design rules of modern tech nodes. [17] is the only ML-based solution that does not require such “non-linear solver” but it as well as other ML-based solutions [18], [15], [16] require tens of thousands of patterns that are the design rule compliant with current settings to train the model. This may not be practical in the early stages of technology node development when DRs are constantly changing, and very few realistic layout patterns, are available. Therefore, to the best of our knowledge, heuristic-based tools continue to lead layout pattern generation in foundries, albeit at a high cost and considerable labor required to adapt to the evolving demands of the foundry.

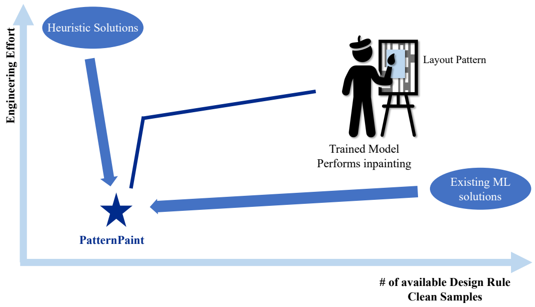

Fig. 1: It is increasingly difficult to find design rule clean training samples (starter patterns) in advanced technology nodes below 10nm due to complex and rigid design rules. Our PatternPaint is designed to work with limited design rule clean samples ($\leq$100) while minimizing the engineering effort.

With recent advancements in pre-trained image foundation models such as DALL-E[21], StatbleDiffusion[22], and Midjourney[23], high-resolution and nearly noise-free images can be directly generated. These models demonstrate significant potential for creating pixel-based layout patterns that accurately conform to specified widths, spacings, and areas, following a straightforward automated denoise process. Moreover, these foundation models also support inpainting[24], redrawing masked regions based on the context of the unmasked areas. For example, an artist may use inpainting to mask a cat and replace it with a dog. In the realm of layout pattern generation, these foundation models can produce DRC (Design Rule Checking)-compliant layout patterns by masking sub-regions of an image and generating variations within these areas. Inspired by these capabilities, this paper introduces **PatternPaint**, an automated, pixel-based layout generation framework that leverages these pre-trained image models. PatternPaint offers significant advantages over existing approaches, as highlighted in Figure 1. Unlike heuristic/rule-based methods, it minimizes engineering efforts and bypasses the need for in-depth design rule knowledge. Compared with existing ML-based methods, PatternPaint does not require training and can generate DRC-compliant images using general pre-trained models, even with limited ($\leq$100) design rule(DR) clean patterns are available. By operating at the pixel level, PatternPaint bypasses the need for solver-based legalization, thus is easily portable and can be used as plug-and-play style pattern generation for nearly all the technology nodes with various number of starter patterns given. With its iterative generation scheme, PatternPaint can produce a wide range of DRC-compliant pattern libraries with limited starter patterns.

Our main contributions are outlined as follows:

- Our work is the first work that uses inpainting in layout generation with a pre-trained image foundation model. This framework, PatternPaint, bypasses the nonlinear solver-based legalization and supports layout pattern generation in complex design rule settings where the existing works struggle to generate legal patterns.
- Our work targeted a realistic industrial scenario where only tens of design rule clean patterns are available and are designed for effortless portability across different technology nodes with minimal engineering effort. Without any training, PatternPaint can produce DR Clean patterns in two completely different design rule settings. With fine-tuning the model, the success rate almost doubled, increasing from 6.25% (pre-finetune) to 11.68% (post-finetune).
- We developed a unique automated template-matching denoising scheme to address noise in generated images under the inpainting scenarios.
- We implement a PCA-based representative sample selection scheme and generate diverse patterns through an iterative process.
- **Validation on industrial PDKs:** PatternPaint is validated on Intel 18A Process Design Kit (PDK) and passed rigorous design rule checking. PatternPaint is able to create more than 4000 DRC-compliant pattern libraries with only 20 starter patterns.

The key novelty of PatternPaint over the prior art is its ability to adapt to complex industrial-standard designs and able to perform with limited starter patterns at the same time, which is crucial for modern technology nodes. This framework is systematic with its unique way to denoise and iterative generation scheme to produce a wide range of DRC-compliant pattern libraries. To the best of our knowledge, PatternPaint is the first unified framework for pattern generation that is able to perform on an industrial technology node. As we describe in later sections, the uniqueness of our framework demonstrates a new paradigm of leveraging image foundation model in layout generation and even layout design which will go beyond the capacity of prior arts.

II. PRELIMINARIES

A. Diffusion model

Diffusion models, also known as Denoising Diffusion Probabilistic Models[25], are generative models in machine learning used for creating complex data such as images, audio, and other high-dimensional entities. These models employ a two-step process: forward diffusion and reverse diffusion. They operate based on a Markov chain that incrementally introduces Gaussian noise over T timesteps, mathematically represented as:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1-\beta_t}x_{t-1}; \beta_t\mathbf{I}) \tag{1}$$

$$q(x_1, ..., x_T|x_0) = \prod_{t=1}^{T} q(x_t|x_{t-1}) \tag{2}$$

where x_0 is the original sample image, $x_1, ..., x_T$ is the noise corrupted sample added with Gaussian noise at each timestep $1, ..., T$, β_t also known as diffusion rate controls the variance schedule where $\beta_t\mathbf{I}$ is the covariance matrix. As T increases and a well-defined β_t schedule is followed, the data distribution gradually approximates a Gaussian distribution, which serves as the starting point for the reverse diffusion process.

$$q(x_T) \approx \mathcal{N}(x_t; 0; \mathbf{I}) \tag{3}$$

In the reverse process, new data samples are generated by reversing the forward diffusion, starting from Gaussian noise:

$$p_\theta(x_0) = \int \prod_{t=1}^{T} p_\theta(x_{t-1}|x_t)dx_{1:T} \tag{4}$$

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t); \sum_\theta(x_t, t)) \tag{5}$$

$$p_\theta(x_T) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t); \sum_\theta(x_t, t)) \tag{6}$$

where θ are learnable deep neural network parameters. The training objective is formulated as follows,

$$L = L_T + L_{T-1} + ... + L_0 \tag{7}$$

$$L_T = D_{KL}(q(x_T|x_0)||p_\theta(x_T)) \tag{8}$$

$$L_t = D_{KL}(q(x_t|x_{t+1}, x_0)||p_\theta(x_t|x_{t+1})), t \in [1, T-1] \tag{9}$$

$$L_0 = -log p_\theta(x_0|x_1) \tag{10}$$

where D_{KL} is the KL divergence, and the term $q(x_t|x_{t+1}, x_0)$ is a tractable reverse conditional probability when conditioned on x_0, which can be derived using Equation (2) and a Bayes'rule. Once a diffusion model is trained, the model can then generate new data samples by sampling from the Gaussian distribution and performing a reverse denoising process as described in the equation above. The generated samples after denoising will lie in the distribution of trained data.

Inpainting[24] is the task of filling masked regions of an image with new content. Inpainting with diffusion models follows a similar reverse process but is conditioned on the known parts of an image. This approach leverages the model's ability to generate coherent structures and textures, filling in the missing pieces in a way that is consistent with the surrounding image data.

B. Squish representation

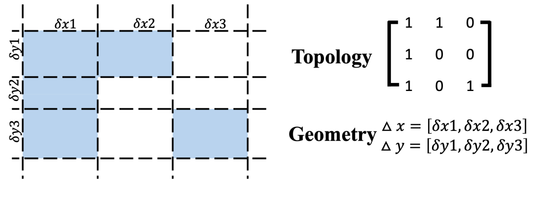

Fig. 2: Squish Pattern Representation.

Standard layout patterns are typically composed of multiple polygons, presenting a sparse informational structure. To efficiently represent these patterns, majority of existing learning-based methods[18],

	Heuristic solutions [13], [14]	DeePattern[15]	CUP [18]	LayouTransformer [17]	DiffPattern[16]	Ours
ML Model Architecture	-	TCAE	VCAE	Transformer	Diffusion model	Pre-trained Diffusion model with Inpainting
Layout representation	-	squish	squish	Sequential Modeling	squish	pixel
Engineering effort	High	Medium	**Low**	**Low**	Medium	**Low**
# of Training samples	None	10k↑	10k↑	10k↑	10k↑	**20 ↓**
Technology-independent	No	No	No	No	No	**Yes**

TABLE I: A Qualitative Comparison with the existing works

[15], [16], [26] use “Squish” pattern[20], [19] to address these issues by compressing a layout into a concise pattern topology matrix alongside geometric data $(\triangle_x, \triangle_y)$, as illustrated in Fig. 2. This process involves segmenting the layout into a grid framework using a series of scan lines that navigate along the edges of the polygon. The distances between each adjacent pair of scan lines are recorded in the $\triangle$ vectors. The topology matrix itself is binary, with each cell designated as either zero (indicating an absence of shape) or one (indicating the presence of a shape).

The Squish representation maintains the physical characteristics of the original patterns through the geometry vector and simplifies the machine learning models’ task to generating topologies. This technique not only compresses layout patterns into a more manageable form but also helps mitigate noise effects in machine-generated images. For example, to represent a rectangle of size (10,20), it is not necessary for all 200 pixels to have a value of 255; instead, a simpler topology of size (2,3) suffices. Current machine learning methods generally reuse existing geometry from their training sets[18] or use a non-linear solver to derive geometry vectors[15], [16]. Each approach has its drawbacks. Using pre-existing geometry vectors limits the diversity of new patterns since no new geometries are introduced. Alternatively, methods incorporating non-linear solvers often face limitations in advanced tech nodes, functioning effectively only in nodes older than 45nm or certain layers of modern tech nodes with unidirectional design rules. Such solvers struggle to adapt as design rules become increasingly complex. To bypass this nonlinear solver, all the prior methods can be reduced by using a pre-defining $\triangle x_i, \triangle y_i$ with a fixed physical width, e.g. each pixel represents a rectangle with 1mm width and 1mm height. We called this method as **pixel-based** solution and prior methods using nonlinear-solver as **squish-based** solutions.

C. Related Works

In the past few years, several synthetic layout pattern generation methods have been proposed. They can be broadly classified into two categories: (a) Rule/Heuristic-based methods, and (b) Machine Learning (ML)-based methods.

Heuristic Methods. The pioneering work in synthetic pattern generation was a heuristic method, proposed in [12]. It used a set of predefined units (small layout building blocks) and applied basic operations such as flipping, rotation, etc., and arranged them randomly to create layout patterns. Later, VIPER [13] introduced a Monte-Carlo approach which employed a constrained random walk algorithm to generate versatile layout patterns. While these methods focused on exploring the design space (i.e., generating diverse patterns), authors of [14] proposed an analytical method to better understand the variations within the known design space (i.e., all possible variations of known layout patterns). Essentially, some polygon edges of existing patterns were varied orthogonally in order to create new patterns with minor local variations. In order to generate DR clean patterns, these heuristic methods require the DRs to be transformed into algorithmic constraints, which need a significant amount of engineering effort not only during the initial development but also while porting to more advanced technology nodes.

ML-based Methods. The emergence of machine learning has spurred the development of generative ML-based solutions for layout pattern generation. DeePattern [15] pioneered this field by using a Variational Autoencoder (VAE)[27] model to generate 1D layout patterns for 7nm EUV unidirectional settings with fixed metal tracks. This method employed a squish representation and a non-linear solver to solve for the geometry vectors. Building on this, CUP[18] expanded the approach to 2D pattern generation, creating a large 2D academic layout pattern dataset of tens of thousands of training samples for a simple design rule setting (minimum width, spacing, and area). They introduced LegalGAN, a generative adversarial network[28], to remove the majority of noises of the generated images. Under this dataset, there are two follow-up works. LayouTransformer first brings transformer techniques in the realm of layout pattern generation and proposes sequential modeling to represent the data. They slightly outperform CUP and show a better diversity score. Later, the existing SOTA DiffPattern[16] proposes to use discrete diffusion methods on layout pattern generation. It came up with a non-linear solver and claimed a 100% success rate with the help of its diffusion model and the solver.

Besides standard layout generation, the additional explorations include CUP-EUV[26] improves upon CUP[18] and demonstrates transferability to reuse trained knowledge from the 2D layout settings to transfer it to EUV-7nm unidirection settings.

Despite its oversimplified design rule settings, existing methods have successfully enriched layout pattern libraries, provided there are enough DR Clean starter samples for training. However, the acquisition of the initial tens of thousands of starter training patterns remains unclear. At the advanced technology nodes below 10nm, this will be hard due to its rigid design rule settings. Additionally, the nonlinear solver-based “legalization” will fail under the realistic complex design rule settings, a problem that we will quantitatively demonstrate in the later of this section. In response, we propose PatternPaint, a pixel-based solution that requires far fewer starter samples. Table I summarizes its advantages over existing methods.

$$
\begin{aligned}
&\delta x_i, \delta y_j > 0, \quad \forall i, j; \\
&\sum_{i=1}^{n} \delta x_i = W, \quad , \sum_{j=1}^{m} \delta y_j = H, \quad \text{(total widths and heights)} \\
&\sum_{i=a}^{b} \delta x_i \geq \text{Space}_{\min}, \quad \forall (a, b) \in \text{SetS}; \\
&\sum_{j=c}^{d} \delta y_j \geq \text{Width}_{\min}, \quad \forall (c, d) \in \text{SetW}; \\
&\delta x_i \cdot \delta y_j \in [\text{Area}_{\min}, \text{Area}_{\max}], \quad \forall \text{ polygon regions}
\end{aligned}
$$

Fig. 5: Nonlinear solver-based legalization [16]

Non-linear solver. Figure 5 presents a non-linear solver that used in prior work[16] as a ”white box” that calculates legal topologies under straightforward design rule settings, such as minimum width, height, and area constraints. However, replicating these methods in more realistic settings involves complexities beyond simple calculations of $\triangle x / \triangle y$, which generally represent polygon widths. Similar to the difference between Side-to-Side (S2S) and Edge-to-Edge (E2E) in Figure 6, the primary direction of a wire must be considered.

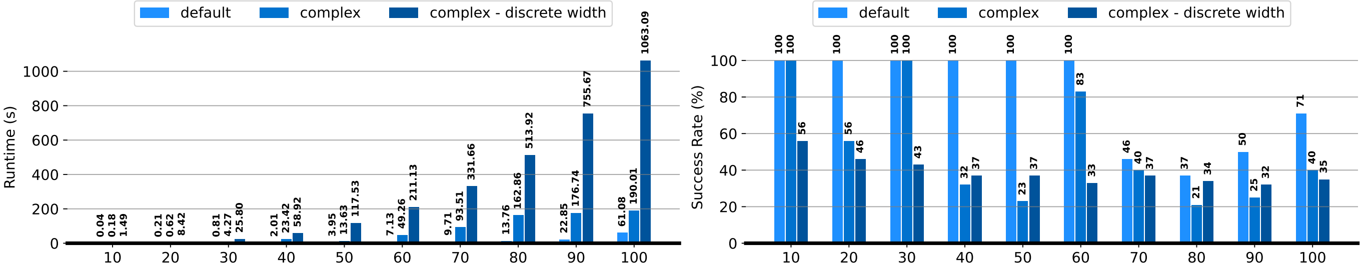

Fig. 3: Runtime and Legality of non-linear solver as topology size increases and complex design rules involves

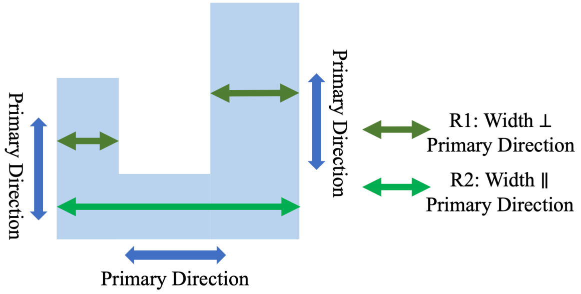

Fig. 4: Realistic Scenario for checking line width

If a wire's horizontal dimension surpasses its vertical, its primary direction is deemed horizontal. Rule 1 (R1) depicted in Figure 4 establishes a basic check for widths that are perpendicular to this primary direction. Furthermore, Rule 2 (R2) evaluates widths that align parallel to the primary direction. Typically, the allowed width range will often involve upper bound and this range will usually be different along horizontal or vertical directions. Extracting SetW from a squish representation to match each specific width range becomes increasingly complex given that we do not know all the shapes beforehand, necessitating extensive validation and substantial engineering effort. At advanced technological nodes, permissible width ranges are often limited to a set of discrete values, significantly complicating the transformation of non-linear programming challenges into mixed integer linear programming problems.

As DiffPattern[16] is not open-sourced and reveals very limited information about their design rule formulations and solver-related implementations, we develop a non-linear solver using the scipy package, specifically employing SLSQP as the solver. We programmed a subset of the design rules commonly used in our design rule runsets, and our findings are detailed in Figure 3. This figure showcases our analysis of the non-linear solver across various topology sizes, where we generated 100 topologies at each size and recorded their average runtime and success rate. For this experiment, we employed three different design rule (DR) settings. The "default" setting adheres to the definitions provided by [16], only checking the minimum width and spacing along with minimum and maximum area constraints. The "complex" setting builds on the "basic", where each width is checked differently in horizontal and vertical directions and includes checks for their minimum and maximum values. Similarly, for spacings, the evaluation considers both horizontal and vertical directions. The "complex-discrete" setting further complicates matters by converting allowed widths into a set of discrete values and only assessing widths perpendicular to the primary direction.

From the "default" implementation to "complex" and then "complex-discrete width", there is a clear increase in runtime, particularly for the "complex-discrete width" configuration due to the inclusion of discrete values. We also note an exponential increase in runtime as topology sizes increase. Success rates, also depicted in Figure 3, indicate that despite all topologies theoretically having a legal solution, the solver consistently struggles with complex rules and the requirement for discrete width values. When dealing with small topology sizes, only the "complex-discrete width" configuration struggles to find legal solutions. However, as topology size expands beyond 60 x 60, all three settings — "default," "complex," and "complex-discrete width" — fail to find more than half of the possible solutions, even though valid solutions exist.

These findings underscore the limitations of the nonlinear solver-based approach in handling larger and more complex topologies. Consequently, a pixel-based method that bypasses these constraints would be more suitable for a broader set of design rule settings, offering a viable alternative for efficient layout generation.

D. Problem Formulation

In this section, we delve into the core principles and evaluation criteria for pattern generation. The objective of pattern generation is to produce a wide variety of realistic layout patterns derived from a small collection of existing designs. A crucial aspect of this process is to ensure that the generated patterns adhere to the established design rules pertinent to actual Integrated Circuit (IC) layouts. These design rules encompass specific geometric parameters such as Space, Width, Area, etc. *Space* is defined as the minimum distance required between two adjacent polygons. It also includes End-to-End (i.e., E2E: minimum distance between edges on the same track) and Side-to-Side (i.e., S2S: minimum distance required between tracks) spacing. *Width* refers to the dimension of a shape along one axis, and *Area* denotes the total surface area enclosed by a polygon. For a layout pattern to be deemed acceptable or 'legal', it must conform to these geometric parameters, which are set at critical threshold values by the design rules.

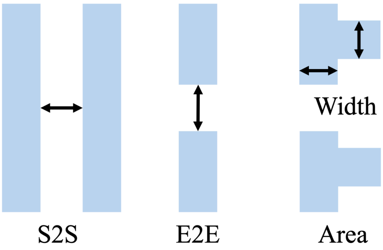

Fig. 6: Illustration of design rules

The diversity and quality of generated patterns are crucial metrics for evaluating the success of layout pattern generation tasks. We examine a variety of metrics to quantitatively analyze the generated patterns. (1) Legality: A layout pattern is legal if and only if it is DR Clean. (2) Entropy – H_1 / $H_{topology}$: As detailed in [16], the complexity of a layout pattern is quantified as a tuple (C_x, C_y) representing the count of scan lines along the x-axis and y-axis, respectively, each reduced by one. Then, we can obtain

$$H_1 = \sum_{i,j} P(C_{x_i}, C_{y_j}) log P(C_{x_i}, C_{y_j})$$

where $P(C_{x_i}, C_{y_j})$ is the probability of encountering a pattern with complexities C_{x_i} and C_{y_j} within the library. Since this metric only focuses on topology diversity without considering any geometric information from actual patterns, we also denote this as $H_{topology}$.

(3) Entropy – H_2 / H_{geo} : To further examine the diversity of actual patterns with their geometric information included, we introduce H_{geo}. For each unique combination of $\triangle x$ and $\triangle y$ (defined in Section II-B) presented in the library, we record their probability $P(\Delta_{x_i}, \Delta_{y_j})$ of having a pattern with the same $\Delta_{x,y}$ matrix within the library.

$$H_2 = \sum_{i,j} P(\Delta_{x_i}, \Delta_{y_j}) log P(\Delta_{x_i}, \Delta_{y_j})$$

(4) Silhouette Score (sil_score) : Silhouette Score is a metric primarily used to measure how well data points are clustered. In our case, we use it to measure the sparsity of layout patterns within and across clusters. It is calculated by

$$Sil\_score = \frac{b-1}{max(a,b)}$$

where a is the mean distance between a sample and all other points in the same class, and b is the mean distance between a sample and all other points in the next nearest cluster. A score close to 1 means that samples are far away from their neighboring clusters and a score close to 0 indicates that the sample is on or very close to the decision boundary between two neighboring clusters. Therefore, in our cases, a lower score indicates higher diversity.

Based on the aforementioned evaluation metrics, the pattern generation problem can be formulated as follows.

Problem 1 (Pattern Generation). *Given a set of design rules and existing patterns, the objective of pattern generation is to synthesize a legal pattern library such that the diversity of the layout patterns in the library is maximized.*

III. FRAMEWORK

A. Overview of PatternPaint

Figure 7 shows the overall framework of PatternPaint which is structured into three primary phases: (1) Pattern Variations Generation. The starter patterns and mask images are input into the pre-trained diffusion model. Noise is randomly sampled in the masked region, and the model performs reverse de-noising process to fill the masked region and create new layout patterns. Each combination of mask and starter image produces multiple new images. (2) Template-based denoising and Design Rule Checking. The generated image is stitched back to the target size using a template-based matching method to mitigate noise effects and remove non-polygon shapes produced by the model. Legal DR clean layouts are collected and DR violated layout patterns are flagged. (3) Iterative Generation. The generated DR Clean Patterns are then collected and perform an iterative generation process. Using a PCA-based selection scheme, we iteratively apply inpainting to create more variations of the generated patterns and generate a diverse pattern library.

In addition to these core components, PatternPaint incorporates a Design Rule Violation Correction feature. This mechanism uses the violated region (output by the DR Checker) as the inpainting mask. This feature efficiently reuses the generated samples and preserves variations that may otherwise be lost in a single inpainting round. Furthermore, PatternPaint is designed to be self-improving. Following the iterative generation and the accumulation of numerous DR clean samples, these patterns are utilized to fine-tune the designated pre-trained model, thus enhancing the efficiency of future generations. The effectiveness of this fine-tuning process in improving generation efficiency will be discussed in a subsequent section V.

B. Pattern Variations Generation

We first prepare starter images and masks as later ML model inputs. Starter Patterns are DR clean layout patterns before we apply PatternPaint to generate layout patterns, for other existing ML approaches, this is also referred as training samples. Black and white image is used as the mask image where white pixels are inpainted and black pixels remain intact. Masks are generated randomly or designated to specific areas where we want to create variations.

Inpainting. Inpainting is then performed by the ML model. During inpainting, a masked image x_0^{masked} is available and the missing parts must be predicted. We can use the forward process to add noise to the known parts, but in the reverse process, we condition on the known pixels. The reverse process for inpainting is slightly modified as:

$$p_\theta(x_{t-1}|x_t, x_0^{masked}) = N(x_{t-1}; \mu_\theta(x_t, x_0^{masked}, t), \sum_\theta(x_t, x_0^{masked}, t)) \tag{11}$$

The mean and covariance now also depend on the original masked image x_0^{masked}, conditioning the reverse process on the known pixels. We also follow the inference scheme mentioned in [22] that only generates masks with about 25% region of its target image size.

Finetune pre-trained model Our framework also supports fine-tuning a pre-trained model to boost the model performance. This fine-tuning shares a similar training objective with Equation (7) but starts from the θ' in the pre-trained neural network. The training objectives then become

$$\begin{aligned} L_{VLB} &= D_{KL}(q(x_T|x_0)||p_{\theta'}(x_T)) - log p_{\theta'}(x_0|x_1) \\ &\quad + \sum_{t=2}^{T} D_{KL}(q(x_t|x_{t+1}, x_0)||p_{\theta'}(x_t|x_{t+1})) \end{aligned} \tag{12}$$

where x_0 is selected from the starter patterns that are available at the fine-tuning stage.

C. Template-based Denoising and DR Checking

Algorithm 1 Template-Matching Denoising

```
Input: Generated image (Noisy), Input image(noise free), Threshold T
Output: Denoise is minimized after this process
Extract scan lines from the noise-free image
Extract scan lines from the noisy image
Cluster similar scan lines within a predefined threshold
for each cluster of squish lines
    Compare clustered lines to scan lines from original image
    if a matching parent scan line is found within the cluster range
        Select the parent scan line for replacement
    else
        Randomly select a line from the cluster for replacement
Apply selected scan lines extract topology matrix from noisy image
Using topology matrix and the selected scan lines, reconstruct the image as denoised Image
return Denoised Image
```

Despite minimal noise in high-resolution images generated by pre-trained models, noise along polygon edges can significantly alter pattern dimensions and easily lead to DR violations. Such noise

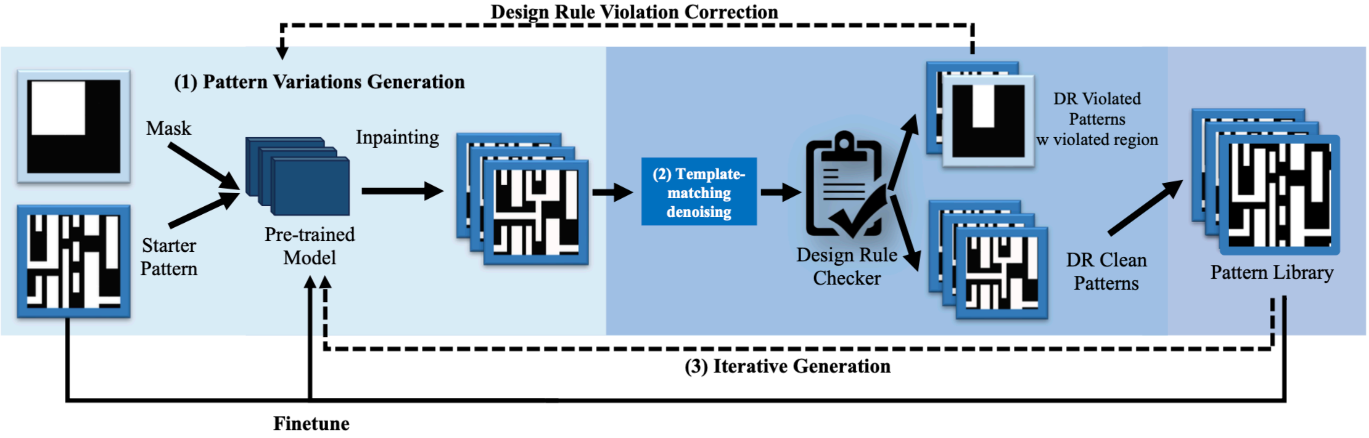
Design Rule Violation Correction
(1) Pattern Variations Generation
Mask
Starter Pattern
Pre-trained Model
Inpainting
(2) Template-matching denoising
Design Rule Checker
DR Violated Patterns w violated region
DR Clean Patterns
Pattern Library
(3) Iterative Generation
Finetune

Fig. 7: Framework of PatternPaint

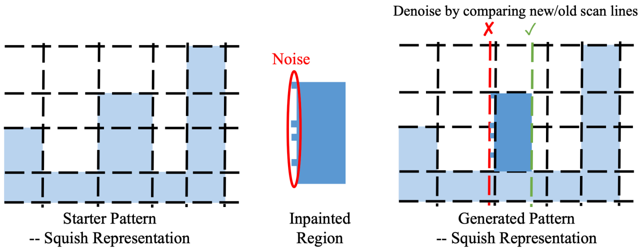

Fig. 8: Illstration of Template-Matching Denoising. Noise at the edge can be reduced by comparing new scan lines with the original scan lines(black). Here, green scan line is preserved since it is larger than a preset threshold, and red scan line is removed.

cannot be avoided since model is based on latent diffusion. In the latent space, a model is lossy and, therefore, leads to information loss in the generated image and cannot perfectly imitate the starter image (i.e., without any noise pixels).

We propose automated template-matching denosing, listed in Algorithm 1, inspired by the fact that only a sub-region of an image is changed during inpainting and the scan lines of the starter pattern (pre-inpainting) are known. We use the squish representation mentioned in Section 2.2, where we first extract scan lines from the noisy generated pattern (post-inpainting) and cluster similar lines within a predefined threshold. We then compare them to scan lines from the template (starter pattern). For each cluster, a parent scan line is chosen if available; otherwise, a line is randomly selected from within the cluster. This method is very effective, and we observe that it significantly increases the number of patterns passing DR checks. Figure 8 also gives an intuitive example of denoising is performed by neglecting unnecessary scan lines due to edge noise but still preserve the scan lines. Denoising is performed by extracting the topology matrix using the designated scan lines and reconstructing the pattern again.

D. Iterative Generation

To produce more new and diverse layout patterns, iterative generation is employed, altering only a sub-region of the image in each iteration. We propose an iterative pattern generation scheme with a Principal Component Analysis (PCA)-based sample selection method, as described in Algorithm 2. PCA reduction provides a qualitative means to illustrate the diversity of a given layout pattern library [13]. The input samples are DR clean layout clips. We first apply PCA to decompose images into several most representative components. To preserve most of the information in the dataset, we push the PCA to have explained_varaince (0.9), meaning 90% of the explained variance is preserved in the dimension-reduced components. Then, an iterative selection is performed to ensure that diverse samples are extracted from the existing library while meeting density constraints. The constraints can be easily integrated with other requirements such as specific pattern shapes or other interesting features, and perform layout pattern generation in a more controlled setting.

The final iteration generation process then integrates Algorithm 2 to select representative starter patterns from the existing pattern library and keep generating layout patterns until the desired diversity is reached or the sample budget is exceeded. When the iterations are completed, a diverse pattern library within the given DR space can be generated.

Algorithm 2 PCA-based Representative and Sparse Sample Selection

```
Input: $X$ (original dataset of size $n\_samples$), $num\_samples$ (number of samples to select), $Sample\_Constraints$
Output: $selected\_samples$ (array of selected representative and sparse samples)
Dimensionality Reduction via PCA:
Apply PCA to $X$ reducing its dimensionality
Choose the number of components to retain significant variance
Let $X\_pca$ be the PCA-reduced dataset
Initialize:
$selected\_indices \leftarrow$ empty list
$remaining\_indices \leftarrow$ list of all indices in $X\_pca$
Select Initial Sample:
Randomly select an initial index from $remaining\_indices$ and add to $selected\_indices$
Remove the selected index from $remaining\_indices$
Iterative Selection:
for $i \leftarrow 1$ to $num\_samples - 1$
    For each index in $remaining\_indices$, calculate the distance
    from the $i$th sample in $X\_pca$ to all samples in $X\_pca[selected\_indices]$
    $farthest\_index \leftarrow$ index with max minimum distance to $selected\_indices$ that also satisfy the sample constrain
    Add $farthest\_index$ to $selected\_indices$
    Remove $farthest\_index$ from $remaining\_indices$
Retrieve Selected Samples:
$selected\_samples \leftarrow X[selected\_indices]$
return $selected\_samples$
```

TABLE II: Quantitative Comparison on synthetic EUV 7nm dataset. A noise is defined as three of its surrounding pixels having the opposite values with the current pixel. Noise level counts the percentage of these noise pixels in an image.

Method	Starter Patterns	Generated Patterns	Legal Patterns	Noise Level (%)
CUP[18](VCAE)	1000	2000	0	8.76
CUP[18](VCAE+LegalGAN)	1000	2000	0	0.54
DiffPattern[16]	20	2000	0	98.7
DiffPattern[16]	200	2000	13	4.02
DiffPattern[16]	1000	2000	566	0.44
PatternPaint[16]	0	2000	**820**	**0.03**

IV. EXPERIMENT: UNI-7NM

In this experiment, we evaluate our solutions on a relatively simple Design Rules setting that only involves unidirectional wires, following the definition of [15]. In this experiment, we will form rigorous comparisons with existing solutions, and quantitatively show that, at the pixel-level, the performance difference with existing solutions.

A. Experimental Setup

Uni-7nm is a synthetic EUV 7nm unidirectional design rule setting following the similar definition of [15]. For this unidirectional dataset, we will only check the E2E spacing as well as the minimum width of a given wire. In this unidirectional setting, all the vertical tracks are fixed so that only horizontal positions (minimum width and E2E) are being checked. We generate 1000 starter patterns with 128 x 128 pixels using a nonlinear solver. Now, we predefined each delta-x and delta-y as 1 pixel, this image also becomes a topology matrix in squish representation and will be used as the training samples for CUP [18] and Diffpattern [16].

PatternPaint Model settings: In this experiment, we employed a pre-trained model (stabldiffusion1.5-inpaint) for pattern generation, denoted as PatternPaint. During inpainting, a mask that masked 1/8 region will be randomly generated and fed to the model, the model will perform inpainting within the masked region of a given starter sample and create new layout patterns.

Summary of Baseline methods: We compare our methods with a Variational Convolution Auto-encoder˜[27] (VCAE) based solution CUP [18] and diffusion-based solution Diffpattern [16]. Our CUP VCAE and LegalGAN implementation follows the network configurations of [18]. VCAE model is trained with 500 epochs with a learning rate of 0.0002. And LegalGAN model is trained with 200 epochs with a learning rate of 0.002. We also follow the [16] to implement a diffusion-based solution. For each diffusion model, we train 400 epochs with a learning rate of 0.0001.

B. Comparison

Table II presents our initial experiments comparing PatternPaint with baseline methods using training sample (starter pattern) sizes of 20, 200, and 1000. Given the prevalence of noise in pixel-based solutions, we quantitatively assessed the noise levels in the generated images. We defined noise pixels as those where three of its surrounding pixels (up, down, left, right) have opposite values. Noise level is defined as counting the percentage of noise pixels in a given image.

1) Noise Level Analysis: Trained with 1000 starter patterns, CUP(VCAE) constantly produces noisy images with a Noise Level of 8%. After performing LegalGAN, the noise level soon reduces to 0.5%, however, the generated images still remain noisy as shown in CUP-1000 in Figure 9. The DiffPattern method exhibits a significant reduction in noise as the number of training samples increases, dropping from 98.7% to 0.44%. PatternPaint achieves the lowest noise levels, registering just 0.03% before implementing our template-based denoising scheme. This performance indicates PatternPaint's potential to directly generate usable layout patterns at the pixel level.

2) Design Rule Checking: Diffusion-based solutions begin to yield DRC-compliant patterns with a minimum of 200 training samples. With 1000 samples, these methods successfully generate compliant patterns in about 25% of cases. PatternPaint demonstrates superior performance by producing 820 legal patterns from 2000 generated images, benefiting from its remarkably low noise level. CUP (VCAE + LegalGAN) fails to produce any DRC-compliant patterns, due to the noise presence in the generated patterns.

3) Qualitative Evaluation: Figure 9 presents a qualitative comparison of layout patterns generated by PatternPaint and other baseline methods. DiffPattern-1k, DiffPattern-200, and DiffPattern-20 represent the diffusion models trained on 1000, 200, and 20 images, respectively. Similarly, CUP-1000 [18] denotes a model trained on 1000 samples. Notably, PatternPaint consistently produces nearly noise-free images, highlighting its superior performance. In contrast, diffusion models only begin to produce legally compliant images when trained with more than 200 samples. The CUP-1000 (VCAE+LegalGAN) model struggles significantly, failing to accurately capture the skeleton of layout patterns. This results in noisy and non-rectangular polygon shapes being generated. fails to produce legal patterns at the pixel-level.

V. EXPERIMENTAL- INTEL 18A

This section formally discussed PatternPaint usage on the industrial standard DR settings (Intel 18A). All the generated patterns are checked through an industry Design Rule Checker. In the following section, we will thoroughly discuss each component of our framework, fine-tuning, iterative generation as well as other usages.

A. Experimental Setup

Intel 18A follows the design rule settings in the latest INTEL 18A setting. The initial datasets only support 20 starter patterns. Since 20 samples are not enough to train diffusion-based solution and VCAE-based solution as shown in section IV, we tried our best to obtain 1000 images from a commercial tool with a size of 512 x 512 pixels. These images were also used in Table IV to evaluate generated pattern performance with a commercial tool.

PatternPaint model setting: In this experiment, besides using stablediffusion1.5-inpaint (PatternPaint), we included another pre-trained foundation model stablediffusion2-inpaint (PatternPaint-ad) [22]. We also explored fine-tuning the model based on the starter images.

Fine-tuning details: We adhered to the fine-tuning procedure described in DreamBooth [29] to fine-tune the inpainting model with our data. The learning rate is set as 5e-6. All the starter clips, a total of 20, are used to fine-tune the model. Fine-tuning on PatternPaint and PatternPaint-ad models takes less than 10 mins on an NVIDIA A100 GPU. For PatternPaint, we denote its fine-tuned model as PatternPaint-finetune. For the PatternPaint-ad model, we denote its fine-tuned model as PatternPaint-finetune. We record the average processing time for each sample in our machine. One Nvidia A100 GPU is used for variation generations and one Intel(R) Xeon(R) Gold 6336Y CPU@2.40GHz is used for the legalization process. The average time for generating a sample was recorded at 0.81 seconds, and for denoising a sample, it was 0.21 seconds.

Baseline Methods: Similar to the Experiment-Uni 7nm, we conducted comparisons using two models, CUP[18] and DiffPattern[16], both trained on 1000 training samples in squish representation [19], [18]. The topology size for these experiments was set at 128 x 128 pixels. DiffPattern, which employs a non-linear solver-based ”legalization” process, initially supported only three basic design

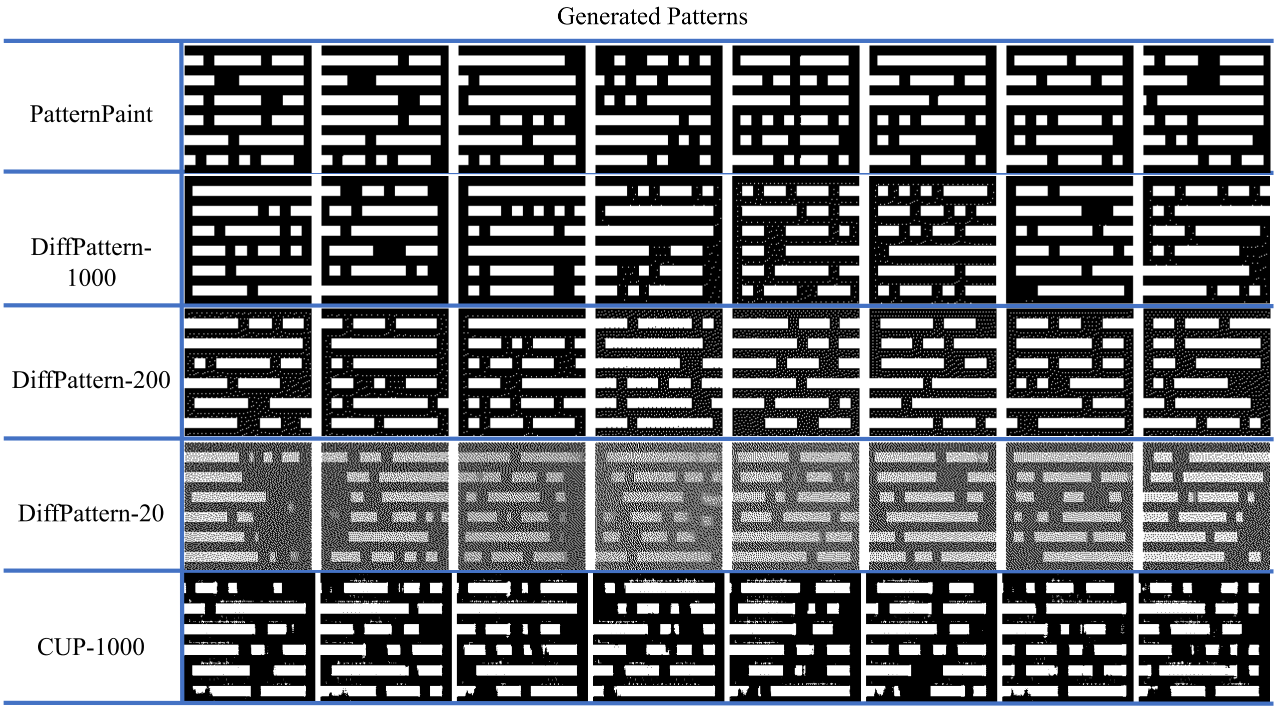

Fig. 9: A Qualitative evaluation of the generated images for CUP [18], DiffPattern [16] as well as Ours work. DiffPattern-1k denoted the diffusion model trained on 1k images, DiffPattern-200 denoted diffusion model trained on 200 images, DiffPattern-20 denoted diffusion model trained on 20 images. Similarly, CUP-1000 [18] is trained 1000 samples. In our settings, only PatternPaint is able to generate nearly noise-free image, Diffusion starts to produce legal images when the starter patterns(for training) are larger than 200.

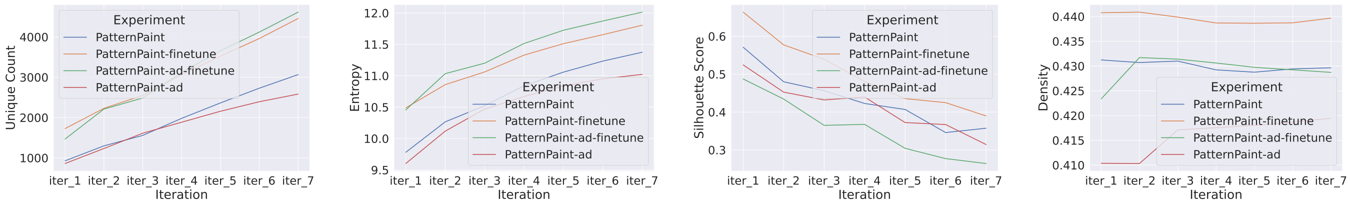

(a) Unique Count (b) Entropy H_2 (c) Silhouette Score (d) Density

Fig. 10: Experiments results for iterative generation process, four metrics are examined: unique pattern counts, the density of generated images, silhouette_score as well as Entropy (H_2)

TABLE III: Runtime comparison with our methods as well as DiffPattern.

Method	Avg Runtime (s)
PatternPaint(Inpainting)	0.81
PatternPaint(Denoising)	0.21
DiffPattern	38.04

rules. However, this approach proved inadequate for our dataset, which includes constraints like discrete values for certain line widths. Despite the engineering challenges, we tried our best to further improve this solver to accommodate a subset of the design rules that involve max-spacing, max-width, and discrete values for certain line widths. After this improvement, legal layout patterns start to appear. We implement this nonlinear solver using scipy package, and the max iteration count is set as 100000000.

TABLE IV: Comparison for pre-trained models as well as internal tool performance

Method	**Generated Patterns**	**Legal Patterns**	**Unique Patterns**	H_1	H_2
Starter Patterns	-	20	20	3.68	4.32
CUP [18]	20000	0	0	0	0
DiffPattern [16]	20000	4	4	2	2
PatternPaint	20000	1251	928	5.06	9.78
PatternPaint-finetune	20000	2336	1728	4.65	**10.49**
PatternPaint-ad	20000	1479	861	5.15	9.60
PatternPaint-ad-finetune	20000	1630	1469	4.96	10.46
Commercial tool	-	1000	1000	**5.70**	9.97

B. Comparison

Table IV showcases our initial experiment using the twenty starter layout clips. We constrained mask selection to ten predefined types to more accurately assess model performance. Diversity metrics H_1 and H_2, calculated by the unique patterns, are also listed. For each combination of clip and mask, our model generated 100 variations, amounting to a total of 20,000 images. These images were subject to DR checks using the entire suite of DRs specified for the given

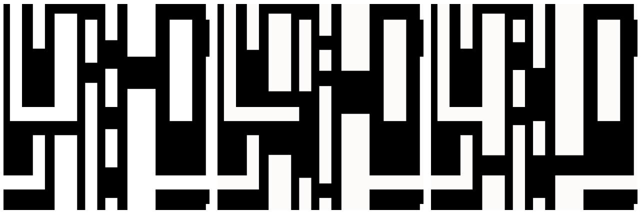

(a) Starter Image (b) Generated Image 1 (c) Generated Image 2

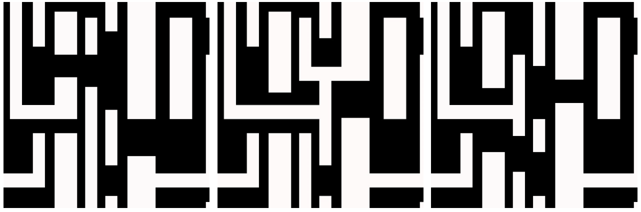

(d) Generated Image 3 (e) Generated Image 4 (f) Generated Image 5

Fig. 11: Generated Variations from a single starter image

layer. Across all four models, we observed the generation of 1000 to 2000 legal patterns. There are a few replicated patterns, however, most of the patterns are unique and they enrich the current pattern library. Notably, there was an observed increase in both H_1 and H_2 for all models compared to the starter patterns. As for the CUP and DiffPattern methods, CUp fails to produce any legal patterns. Diffpattern generates a very limited number of legal patterns (four) after the "legalization" process, which also requires significantly more runtime. This suggests that this solver-based solution may not be well-suited for complex design rule settings. Table III details the average time required to generate and process a layout pattern. Currently, DiffPattern takes approximately 30 times longer than PatternPaint to process a generated layout pattern. Although improving this non-linear solver to accommodate a broader set of design rules might increase the yield of legal patterns, a substantial increase in the runtime will also be expected.

The effectiveness of the proposed fine-tuning process is evident by comparing the PatternPaint model with PatternPaint-finetune, and PatternPaint-ad with PatternPaint-ad-finetune. It is seen by an increase in both the count of legal patterns and the number of unique patterns for the fine-tuned models. Diversity value, calculated using the aforementioned equation and legal patterns, shows that the fine-tuned PatternPaint-finetune model has the best performance. There is a significant increase in diversity compared to both the starter patterns and our model. However, an internally utilized heuristic-based pattern generation tool still surpassed all other methods. This superiority is attributed to the tool's manually encoded technology node and DR information, enabling efficient exploration within permissible ranges. The proposed model, on the other hand, relies on neighborhood information and random steps, and makes regional changes to 20 starter layout patterns, therefore, leading to a more clustered pattern library in the first round. In the later steps, however, we notice that more variations are produced through the iterative generation process. Figure 11 visually represents the variations generated by the proposed methods. While the starter image is depicted in (a), (b-f) show the generated images. We notice that the proposed methods explored a wide range of variations, demonstrating the models' awareness of tracks. For example, in (f), the model attempts to disconnect from an adjacent thick track and establish a connection with a farther one. In (e), more complex changes are made, forming connections with even farther tracks and upper objects.

C. Iterative Pattern Generation

Following the initial experiment, we created a library of unique patterns with substantial variation. We then embarked on iterative pattern generation, as described in section III-D, to check if diversity increased through this process. We designated the unique patterns from experiment 1 (Table 1) as our first iteration. For subsequent iterations, we conducted PCA analysis to select 100 of the most sparse representative samples, with the density constraint set at 40% for the selected patterns. Using five predefined masks, we generated 5000 samples out of the 100 patterns, adding only clean and new samples to our existing pattern library.

Figure 10 shows the experimental results as iterations proceed. We evaluated four metrics: unique pattern count, image density, silhouette score, and entropy (H_2). Both unique pattern count and H_2 increased with each iteration. A notable gap was observed between baseline models (PatternPaint, PatternPaint-ad) and the fine-tuned models (PatternPaint-finetune, PatternPaint-ad-finetune), with the latter consistently outperforming the former. This further validates our fine-tuning process, demonstrating significant model improvements. Then, we examine the density of patterns in the current pattern library. This is because we observed that during iterative pattern generation, some models tend to take the easy path and produce sparse, empty areas in the masked region. While they are still legal and new patterns, these patterns are not useful. Therefore, we also keep track them during this process. For the starter images, the density is around 43.6%. We noticed that only PatternPaint-ad-finetune matches the density in the starter layout clips. For all the other models, slight decreases are observed with PatternPaint and PatternPaint-finetune. For PatternPaint-ad, however, a higher density drop is observed, indicating that some sparse region exists in this existing pattern library.

Lastly, we examine the silhouette score of our existing pattern library at each iteration. In order to obtain sil_score, we first decrease the dimensionality of images using PCA. We use a large number of principal components to preserve 90% of the explained variance in the dataset. Then we perform K-means clustering for these decomposed components. A number of 20 clusters are selected since we start from 20 distinct starter layout clips. An average sil_score is computed for each iteration and shown in Figure 10. As iterations increase, we observe a constant decrease in sil_score, indicating better diversity. PatternPaint-ad-finetune significantly outperforms other three models with lower sil_score. This shows that PatternPaint-ad-finetune is able to produce more interesting variations than other models and the variations are not limited to the space around starter patterns. This also aligns with our observations in density analysis where PatternPaint-ad-finetune is able to produce high-density patterns that involve more risky changes.

D. Violation Correction

To further demonstrate the advantages of PatternPaint and its ability to take in DRC feedback as the violated region, we present an application scenario as a case study below.

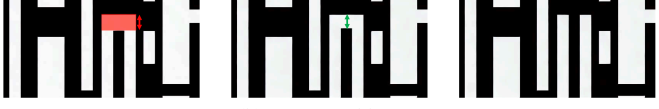

(a) DRV layout (b) Recovered layout with correct width (c) DR Clean layout

Fig. 12: DRV layout and fixed layouts

1) Case Study 1: Single Design Rule Correction: Figure 12 details our experiment on a single Design Rule Violated (DRV) layout, where the DRC flagged a region (highlighted in red) due to incorrect width. We explored different masking options: Mask-0 used the original flagged region, Mask-1 expanded it by 25%, and Mask-2 by 50%. PatternPaint generated 20 variations for each mask, successfully exploring the correct widths as shown in Figure 9(b). Notably, our model identified feasible widths from a limited set within 60 attempts. Figure 12 (c) shows another attempt to remove the connection between neighboring tracks that results in a DR Clean layout.

As shown above, meaningful variations are being generated using the violated region as the inpainting mask. This correction method poses an interesting feature that could go beyond the scope of test pattern generation. Specifically, during the VLSI layout design phase, our framework could be seamlessly integrated with an existing Design Rule Checker to automate the correction of design rule violations. This capability opens up new avenues for a range of applications where prior work fails to explore as they cannot work with specific regional changes.

VI. CONCLUSION

In this paper, we propose PatternPaint, an automated pattern generation framework via inpainting techniques. We implement our own unique template-matching denosing scheme to tackle noise and propose a PCA-based sample selection scheme for iterative pattern generation. In the initial round of the iterative generation process, thousands of DR clean layouts are generated on the latest Intel PDK and checked through an industry-standard DR checker. In later iterations, by measuring entropy and silhouette scores, we observed that pattern diversity improves. Our work, PatternPaint, has its unique benefits with little to no human effort in loop and can work in cases even when starter patterns are highly limited. This framework is also technology node independent, without any prior knowledge of the design rules of the technology node, we demonstrate that PatternPaint can generate various legal patterns across two different design rule settings. In future works, we hope to explore a broader selection of the pre-trained model, further fine-tune the pre-trained foundation models with more layout samples, randomize mask generation, and launch a larger batch of experiments to compare the explored design rule space with real layout patterns from products.

REFERENCES

[1] J.-R. Gao, X. Xu, B. Yu, and D. Z. Pan, "Mosaic: Mask optimizing solution with process window aware inverse correction," in *Proceedings of the 51st Annual Design Automation Conference*, DAC '14, (New York, NY, USA), p. 1–6, Association for Computing Machinery, 2014.

[2] Y. Jiang, F. Yang, B. Yu, D. Zhou, and X. Zeng, "Efficient layout hotspot detection via neural architecture search," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 27, jun 2022.

[3] B. Jiang, H. Zhang, J. Yang, and E. F. Y. Young, "A fast machine learning-based mask printability predictor for opc acceleration," in *Proceedings of the 24th Asia and South Pacific Design Automation Conference*, ASPDAC '19, (New York, NY, USA), p. 412–419, Association for Computing Machinery, 2019.

[4] J. Kuang, W.-K. Chow, and E. F. Y. Young, "A robust approach for process variation aware mask optimization," in *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, DATE '15, (San Jose, CA, USA), p. 1591–1594, EDA Consortium, 2015.

[5] H. Yang, S. Li, Y. Ma, B. Yu, and E. F. Y. Young, "Gan-opc: mask optimization with lithography-guided generative adversarial nets," in *Proceedings of the 55th Annual Design Automation Conference*, DAC '18, (New York, NY, USA), Association for Computing Machinery, 2018.

[6] G. R. Reddy, K. Madkour, and Y. Makris, "Machine learning-based hotspot detection: Fallacies, pitfalls and marching orders," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pp. 1–8, 2019.

[7] R. Chen, W. Zhong, H. Yang, H. Geng, X. Zeng, and B. Yu, "Faster region-based hotspot detection," in *Proceedings of the 56th Annual Design Automation Conference 2019*, DAC '19, (New York, NY, USA), Association for Computing Machinery, 2019.

[8] J. Pan, C.-C. Chang, Z. Xie, J. Hu, and Y. Chen, "Robustify ml-based lithography hotspot detectors," in *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, ICCAD '22, (New York, NY, USA), Association for Computing Machinery, 2022.

[9] H. Zhang, B. Yu, and E. F. Young, "Enabling online learning in lithography hotspot detection with information-theoretic feature optimization," in *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pp. 1–8, 2016.

[10] H. Yang, Y. Lin, B. Yu, and E. F. Y. Young, "Lithography hotspot detection: From shallow to deep learning," in *2017 30th IEEE International System-on-Chip Conference (SOCC)*, pp. 233–238, 2017.

[11] C. Y.-C. Hou, "Design challenges and enablement for 28nm and 20nm technology nodes," in *2010 Symposium on VLSI Technology*, pp. 225–226, 2010.

[12] H. Li, E. Zou, R. Lee, S. Hong, S. Liu, J. Wang, C. Du, R. Zhang, K. Madkour, H. Ali, *et al.*, "Design space exploration for early identification of yield limiting patterns," in *Design-Process-Technology Co-optimization for Manufacturability X*, vol. 9781, 2023.

[13] G. R. Reddy, M.-M. Bidmeshki, and Y. Makris, "Viper: A versatile and intuitive pattern generator for early design space exploration," in *2019 IEEE International Test Conference (ITC)*, pp. 1–7, 2019.

[14] G. R. Reddy, C. Xanthopoulos, and Y. Makris, "Enhanced hotspot detection through synthetic pattern generation and design of experiments," in *2018 IEEE 36th VLSI Test Symposium (VTS)*, pp. 1–6, 2018.

[15] H. Yang, S. Li, W. Chen, P. Pathak, F. Gennari, Y.-C. Lai, and B. Yu, "Deepattern: Layout pattern generation with transforming convolutional auto-encoder," *IEEE Transactions on Semiconductor Manufacturing*, vol. 35, no. 1, pp. 67–77, 2022.

[16] Z. Wang, Y. Shen, W. Zhao, Y. Bai, G. Chen, F. Farnia, and B. Yu, "Diffpattern: Layout pattern generation via discrete diffusion," in *2023 60th ACM/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2023.

[17] L. Wen, Y. Zhu, L. Ye, G. Chen, B. Yu, J. Liu, and C. Xu, "Layoutransformer: Generating layout patterns with transformer via sequential pattern modeling," in *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, ICCAD '22, (New York, NY, USA), Association for Computing Machinery, 2022.

[18] X. Zhang, J. Shiely, and E. F. Young, "Layout pattern generation and legalization with generative learning models," in *2020 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, pp. 1–9, 2020.

[19] H. Yang, P. Pathak, F. Gennari, Y.-C. Lai, and B. Yu, "Detecting multilayer layout hotspots with adaptive squish patterns," in *Proceedings of the 24th Asia and South Pacific Design Automation Conference*, ASPDAC '19, (New York, NY, USA), p. 299–304, Association for Computing Machinery, 2019.

[20] F. E. Gennari and Y.-C. Lai, "Topology design using squish patterns," *U.S. Patent 8832621B1*, 2014.

[21] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," 2022.

[22] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.

[23] I. Midjourney, "Midjourney," 2022.

[24] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "Repaint: Inpainting using denoising diffusion probabilistic models," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11451–11461, 2022.

[25] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *arxiv:2006.11239*, 2020.

[26] X. Zhang, H. Yang, and E. F. Young, "Attentional transfer is all you need: Technology-aware layout pattern generation," in *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pp. 169–174, 2021.

[27] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2022.

[28] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5967–5976, 2017.

[29] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, "Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.