

DeepRWCap: Neural-Guided Random-Walk Capacitance Solver for IC Design

Hector Rodriguez Rodriguez*, Jiechen Huang*, Wenjian Yu

Department of Computer Science and Technology, BNRist, State Key Laboratory of Cryptography and Digital Economy Security, Tsinghua University, Beijing 100084, China
{lad24,hjc22}@mails.tsinghua.edu.cn, yu-wj@tsinghua.edu.cn

Abstract

Monte Carlo random walk methods are widely used in capacitance extraction for their mesh-free formulation and inherent parallelism. However, modern semiconductor technologies with densely packed structures present significant challenges in unbiasedly sampling transition domains in walk steps with multiple high-contrast dielectric materials. We present DeepRWCap, a machine learning-guided random walk solver that predicts the transition quantities required to guide each step of the walk. These include Poisson kernels, gradient kernels, signs and magnitudes of weights. DeepRWCap employs a two-stage neural architecture that decomposes structured outputs into face-wise distributions and spatial kernels on cube faces. It uses 3D convolutional networks to capture volumetric dielectric interactions and 2D depthwise separable convolutions to model localized kernel behavior. The design incorporates grid-based positional encodings and structural design choices informed by cube symmetries to reduce learning redundancy and improve generalization. Trained on 100,000 procedurally generated dielectric configurations, DeepRWCap achieves a mean relative error of $1.24 \pm 0.53\%$ when benchmarked against the commercial Raphael solver on the self-capacitance estimation of 10 industrial designs spanning 12 to 55 nm nodes. Compared to the state-of-the-art stochastic difference method Microwalk, DeepRWCap achieves an average 23% speedup. On complex designs with runtimes over 10 s, it reaches an average 49% acceleration.

Code — <https://github.com/THU-numbda/deepRWCap>

Extended version — <https://arxiv.org/abs/2511.06831>

1 Introduction

The extraction of parasitic capacitance is a key step in the design and verification of Integrated Circuits (ICs), which involves analyzing the physical layout before fabrication to ensure that it meets timing, power consumption, and signal integrity requirements. Accurate capacitance extraction (typically within 5% error) has become increasingly challenging as the semiconductor industry shifts toward more complex process technologies (Yu, Song, and Yang 2021). With the deceleration of Moore’s Law, the industry has shifted from 2D to 3D integration, leveraging tightly packed

architectures to sustain performance improvements. Transistor structures evolve from FinFETs to Gate-All-Around and Complementary Field-Effect Transistors (Mukesh and Zhang 2022; Liebmann et al. 2021). At the circuit level, vertically stacked ICs enable smaller footprints while reducing interconnect lengths (Gomes et al. 2022; Wu et al. 2022).

These increasingly complex 3D structures pose significant computational challenges for capacitance extraction. Traditional methods for capacitance extraction struggle to balance computational accuracy and efficiency in modern IC design. Electrostatic field solvers such as the finite-difference method (FDM) deliver high accuracy but suffer from poor scalability when applied to large structures (Yu et al. 2013). Pattern matching-based approaches offer computational efficiency and scalability, but provide no accuracy guarantees and rely heavily on specialized domain expertise (Yang et al. 2023).

Random walk methods provide a scalable stochastic framework for solving elliptic partial differential equations (PDEs) such as Poisson and diffusion equations. Originating from the classical Walk-on-Spheres (WoS) algorithm (Muller 1956), they leverage the duality between PDEs and stochastic processes. They have gained renewed interest in computer graphics and deep learning applications (Sawhney and Crane 2020; Li et al. 2023; Miller et al. 2024; Nam, Berner, and Anandkumar 2024) due to the inherent parallelism and controllable accuracy. However, while WoS has been extended to support varying diffusion coefficients (Sawhney et al. 2022), it requires the coefficient to be twice-differentiable and is inapplicable to the piecewise constant dielectrics in ICs. Thus, the random walk method for capacitance extraction still requires computationally expensive transition evaluations to guide each step when dealing with complex multi-dielectric domains (Visvardis, Liaskovitis, and Efstathiou 2023; Huang and Yu 2024, 2025).

Recent advances in machine learning have introduced AI-based techniques across various Electronic Design Automation (EDA) applications, including code generation, placement, and routing optimization (Cheng et al. 2022; Zhong et al. 2024; Lai et al. 2025). In the domain of capacitance extraction, learning-based methods either replace traditional solvers by directly predicting capacitance from layout geometry (Yang et al. 2023; Liu et al. 2024; Cai et al. 2024) or assist core computational steps in numerical algorithms such

*These authors contributed equally.

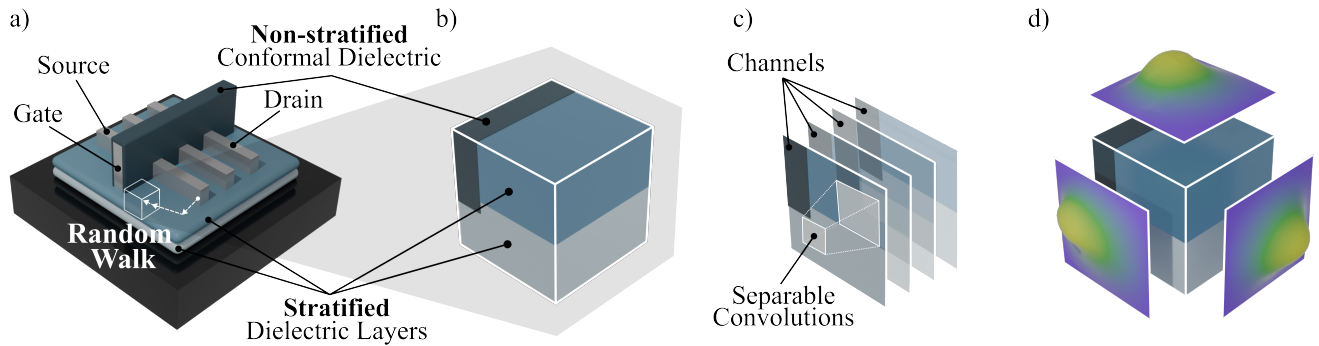


Figure 1: Illustration of the neural-guided random walk capacitance extraction framework: (a) FinFET structure, (b) transition cube, (c) compact CNN, and (d) predicted kernels.

as random walk methods (Visvardis, Liaskovitis, and Efstathiou 2023). However, the fully learning-based methods struggle to generalize across different designs and technology nodes, and the hybrid methods do not offer a compelling efficiency-accuracy trade-off.

In this work, we present DeepRWCap, which builds upon a neural-guided random walk framework as illustrated in Fig. 1. The framework leverages compact CNNs to accelerate the random walk-based capacitance extraction. In Fig. 1, starting from a FinFET structure embedded in heterogeneous dielectrics (Fig. 1a), a local transition cube (Fig. 1b) is extracted to represent the dielectric environment surrounding each walk step. This cube is processed by compact, depthwise-separable CNNs (Fig. 1c) that estimate one of the transition quantities, the Poisson kernel (Fig. 1d) to inform the next move in the walk. Our contributions are threefold: (1) a unified neural architecture for predicting both Poisson and gradient kernels across multi-dielectric domains, (2) a GPU-accelerated inference engine with producer-consumer scheduling for high-throughput sampling, and (3) an extensive evaluation on 10 industrial test cases with different technology nodes, showing an average 23% speedup over the state-of-the-art method (Huang and Yu 2025) while preserving high accuracy with a mean relative error of 1.24%.

2 Related Work

Monte Carlo Methods for Capacitance Extraction The Floating Random Walk (FRW) method (Y. L. Le Coz and R. B. Iverson 1992; Yu et al. 2013) adapts Monte Carlo techniques to rectilinear IC geometries using cubic transition domains. Several approaches address the computational challenge of transition kernel evaluation: OCT (Yang and Yu 2020) approximates dielectric configurations through volume-weighted averaging but introduces errors in high-contrast scenarios; AGF (Huang and Yu 2024) provides exact solutions for stratified dielectrics but requires cube shrinking or volume-weighted averaging to handle non-stratified cases, increasing the Monte Carlo variance or introducing approximation errors; Microwalk (Huang and Yu 2025) achieves unbiased FRW transitions but lacks support for the first step of each walk (gradient sampling), requiring fallbacks to AGF and cube shrinking.

Learning-Based Methods for Capacitance Extraction. CNN-Cap (Yang et al. 2023) employs ResNet architectures to directly predict layout capacitances by representing conductors as binary masks across different layers. GNN-Cap (Liu et al. 2024) uses Graph Convolutional Networks with spatial message passing to extract capacitances, achieving comparable accuracy with better scalability than CNN-based approaches. While these methods offer speedups over electrostatic solvers, they require retraining for each process node since dielectric distributions are learned implicitly rather than explicitly modeled. PCT-Cap (Cai et al. 2024) uses Point Cloud Transformers with 8-dimensional feature vectors that include spatial coordinates, normal vectors, and relative dielectric permittivities. While achieving higher accuracy than CNN-based methods, it approximates multi-dielectric scenarios through local averaging around each point, which inadequately captures the global dielectric structure between conductors, limiting its applicability to conformal dielectrics.

Hybrid Methods for Capacitance Extraction. Visvardis et al. present deep learning-driven random walks using Group-Equivariant Convolutional Neural Networks (GECNN) (Worrall and Brostow 2018) with autoencoder compression and Gaussian Mixture Models (GMM) for sampling (Visvardis, Liaskovitis, and Efstathiou 2023). However, several limitations hinder practical deployment: the group equivariant representation creates information bottlenecks by aggressively reducing the spatial dimensions, the synchronous walker architecture underutilizes parallelism, and the CPU-only implementation results in runtimes that are $12\times$ slower than vanilla FRW methods (Visvardis, Liaskovitis, and Efstathiou 2023). Additionally, the restricted dielectric range $[1, 10]$ and GMM’s inability to capture sharp discontinuities from high-contrast interfaces limit applicability to advanced semiconductor technologies.

3 Preliminaries

3.1 Capacitance Extraction

For a system of N_c conductors in a three-dimensional, multi-dielectric domain $\Omega \subset \mathbb{R}^3$, the *capacitance matrix* $\mathbf{C} \in \mathbb{R}^{N_c \times N_c}$ describes the linear relation between the

electric potentials $\mathbf{u} \in \mathbb{R}^{N_c}$ and the conductor charges $\mathbf{q} \in \mathbb{R}^{N_c}$ (Smolić and Klajn 2021):

$$\mathbf{q} = \mathbf{C}\mathbf{u}. \quad (1)$$

The diagonal element C_{ii} is the *self-capacitance* of conductor i , and C_{ij} for $i \neq j$ is the *coupling capacitance* between conductors i and j .

To compute C_{ij} , we set up an electrostatic problem where conductor j is held at a unit potential ($u_j = 1$) and all other conductors are grounded ($u_k = 0, \forall k \neq j$). The target capacitance is then equal to the charge on conductor i , which can be expressed via Gauss’s law (Griffiths 2023) as

$$C_{ij} = q_i = - \oint_{G_i} \alpha(\mathbf{r}) \frac{\partial u(\mathbf{r})}{\partial \mathbf{n}_r} ds, \quad (2)$$

where G_i is a Gaussian surface enclosing conductor i , α is the spatially varying dielectric permittivity (i.e., the diffusion coefficient in Laplace’s equation), and \mathbf{n}_r is the unit outward normal vector at \mathbf{r} . The electrostatic field $u(\mathbf{r})$ satisfies the following boundary value problem of Laplace’s equation:

$$\begin{cases} \nabla \cdot (\alpha(\mathbf{r}) \nabla u(\mathbf{r})) = 0, & \mathbf{r} \in \Omega, \\ \frac{\partial u(\mathbf{r})}{\partial \mathbf{n}_r} = 0, & \mathbf{r} \in \partial\Omega_0, \\ u(\mathbf{r}) = u_k, & \mathbf{r} \in \partial\Omega_k, 1 \leq k \leq N_c, \end{cases} \quad (3)$$

where the Neumann boundary $\partial\Omega_0$ represents the extraction window and the Dirichlet boundaries $\partial\Omega_k$ are the conductor surfaces, such that $\cup_{k=0}^{N_c} \partial\Omega_k = \partial\Omega$ and $\partial\Omega_{k_1} \cap \partial\Omega_{k_2} = \emptyset, \forall k_1 \neq k_2$. In essence, capacitance extraction is applying numerical methods to solve (3) and evaluate the surface integral in (2).

3.2 Floating Random Walk Method

The connection between PDEs and stochastic processes is well established through Kakutani’s theorem (Kakutani 1944) and, more generally, Feynman-Kac formula (Øksendal 2003). Specifically, the solution to Laplace’s equation has the following stochastic representation

$$u(\mathbf{r}) = \mathbb{E}[u(\mathbf{W}_\tau^\xi) | \xi = \mathbf{r}], \quad (4)$$

where \mathbf{W}_t^ξ is a standard Brownian motion starting at ξ , and τ is the first hitting time of \mathbf{W}_t^ξ to the Dirichlet boundary.

To adapt to the rectilinear geometries of circuit layouts, the floating random walk (FRW) method uses cubes to “hop” in the problem domain (Y. L. Le Coz and R. B. Iverson 1992). The resulting FRW process forms a discrete-time Markov chain \mathbf{X}_k , where each transition is sampled on the surface of the largest cube centered at \mathbf{X}_k (the so-called *transition cube*). The transition probability kernel is the *Poisson kernel* p_α (Axler, Bourdon, and Wade 2001)¹. It is the normal derivative of the Green’s function for Laplace’s

¹It is also referred to as “surface Green’s function” in literature (Visvardis, Liaskovitis, and Efstathiou 2023; Huang and Yu 2024).

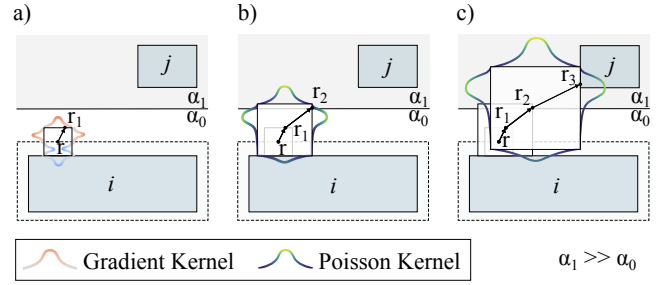


Figure 2: An FRW trajectory from conductor i to j with example transition kernels, showing (a) the initial step, (b) an intermediate step, and (c) the boundary hit.

equation in the cube and depends on the dielectric configuration α within the cube. Based on the FRW process, the solution to (3) is expressed as

$$\begin{aligned} u(\mathbf{r}_k) &= \mathbb{E}[u(\mathbf{X}_{k+1}) | \mathbf{X}_k = \mathbf{r}_k] \\ &= \oint_{\partial S(\mathbf{r}_k)} u(\mathbf{r}_{k+1}) p_\alpha(\mathbf{r}_{k+1} | \mathbf{r}_k) ds, \end{aligned} \quad (5)$$

where $S(\mathbf{r}_k)$ is the transition cube centered at \mathbf{r}_k . We recursively apply Monte Carlo integration based on (5) until hitting the Dirichlet boundary at some \mathbf{r}_T , in which case the boundary condition $u(\mathbf{r}_T)$ is an unbiased estimator for $u(\mathbf{r}_k)$. Combining this estimator and our target integral in (2), we obtain the FRW estimator for capacitance

$$\hat{q}_i = - \frac{\alpha(\mathbf{r})}{f(\mathbf{r})} w_\alpha(\mathbf{r}) s_\alpha(\mathbf{r}, \mathbf{r}_1) u(\mathbf{r}_T), \quad (6)$$

where $f(\mathbf{r})$ is the probability density on G_i , $w_\alpha(\mathbf{r}) = \oint_{\partial S(\mathbf{r})} \left| \frac{\partial p_\alpha(\mathbf{r}_1 | \mathbf{r})}{\partial \mathbf{n}_r} \right| ds_1$, and $s_\alpha(\mathbf{r}, \mathbf{r}_1) = \text{sign}\left(\frac{\partial p_\alpha(\mathbf{r}_1 | \mathbf{r})}{\partial \mathbf{n}_r}\right)$. The transition kernel is $g_\alpha(\mathbf{r}_1 | \mathbf{r}) = \left| \frac{\partial p_\alpha(\mathbf{r}_1 | \mathbf{r})}{\partial \mathbf{n}_r} \right| / w_\alpha(\mathbf{r})$ for the first transition and $p_\alpha(\mathbf{r}_{k+1} | \mathbf{r}_k)$ for the subsequent transitions (Yu et al. 2013). See (Huang, Yang, and Yu 2024) for the derivation and proof.

4 DeepRWCap: Neural-Guided Random Walk Solver

DeepRWCap is a neural-guided framework that accelerates random walk-based capacitance extraction by learning to predict transition kernels in non-stratified domains. We formulate this as a supervised learning problem and describe our synthetic dataset generation approach for training on procedurally generated dielectric configurations. We present a neural architecture that exploits cube symmetries to avoid learning redundancy and integrates with a high-throughput inference engine for efficient random walk sampling.

4.1 Problem Formulation: Multi-Dielectric Transition Kernel Learning

For a transition cube $S \subset \mathbb{R}^3$, let $\alpha : S \rightarrow \mathbb{R}_{>0}$ denote the relative permittivity, which is a piecewise constant function in actual IC technologies, representing the planar dielectric layers and layout-dependent dielectric blocks (see Fig.

1). Performing an unbiased FRW transition requires a set of transition quantities: the *Poisson kernel* $p_\alpha : \partial S \rightarrow \mathbb{R}_{>0}$, the *weight value* $w_\alpha \in \mathbb{R}_{>0}$, the *sign distribution* $s_\alpha : \partial S \rightarrow \{-1, 0, 1\}$, and the *gradient kernel* $g_\alpha : \partial S \rightarrow \mathbb{R}_{\geq 0}$. Note that they depend only on the relative dielectric configuration inside S and are independent of the absolute location of the cube. Therefore, the learning problem is to approximate the mapping from α to $w_\alpha, s_\alpha, g_\alpha$ for the first transition, and p_α for the subsequent transitions.

We discretize the transition cube on an $N \times N \times N$ voxel grid, where the dielectric α is represented as a tensor $\mathcal{X} \in \mathbb{R}^{N \times N \times N}$. Since the target quantities are invariant to global scaling of α , we normalize each input such that $\mathcal{X}_{i,j,k} = \frac{\alpha_{i,j,k}}{\max(\alpha)} \in (0, 1]$. The target surface functions are discretized over the six cube faces as tensors $\mathbf{s}_\alpha, \mathbf{g}_\alpha, \mathbf{p}_\alpha \in \mathbb{R}^{6 \times N \times N}$. This not only provides a robust and high-fidelity data representation, but also aligns with the standard finite-difference scheme in numerical solutions of the Poisson kernel (Huang and Yu 2025).

4.2 Training Dataset Generation

The FRW transition quantities exhibit a highly non-linear dependence on the dielectric configuration (Huang and Yu 2024) and the governing equations are not available in explicit forms. This precludes the use of physics-informed objectives (Raissi, Perdikaris, and Karniadakis 2019) and motivates a fully data-driven learning framework.

To enable models that generalize across different semiconductor technologies, we synthesize a dataset $\mathcal{D} = \{w_\alpha^{(i)}, \mathbf{s}_\alpha^{(i)}, \mathbf{g}_\alpha^{(i)}, \mathbf{p}_\alpha^{(i)}; \mathcal{X}^{(i)}\}$ using a block-based generative procedure (Alg. 1) that emulates real-world dielectric patterns. Each block is parameterized by its center (c_x, c_y, c_z) , side lengths (l_x, l_y, l_z) , and relative permittivity κ . To mimic materials in practical technologies, we empirically design RandomDielectric in Alg. 1 as sampling a low- κ distribution $U(2, 10)$ with 80% probability or a high- κ distribution $U(10, 80)$ with 20% probability. Nested structures are randomly added (line 6-11 in Alg. 1) to mimic conformal coatings. Blocks in the list can overlap, with earlier ones overriding later ones. This algorithm produces high-contrast overlaps and sharper kernels compared to real layout dielectrics.

After placing the blocks via Alg. 1, the domain $[-1, 1]^3$ is treated as the transition cube and voxelized as the tensor representation $\mathcal{X}^{(i)}$. With a given $\mathcal{X}^{(i)}$, we use a finite-difference-method (FDM) solver (Huang and Yu 2025) to obtain the targets $w_\alpha^{(i)}, \mathbf{s}_\alpha^{(i)}, \mathbf{g}_\alpha^{(i)}, \mathbf{p}_\alpha^{(i)}$.

4.3 Network Architecture for Transition Kernel Prediction

The Poisson kernel is equivariant to the symmetry group of the cube, meaning that for any symmetry (including rotations, reflections, and inversions), the kernel values permute and transform consistently with the geometric transformation of the cube.

Thus, predicting this $\mathbf{p}_\alpha \in \mathbb{R}^{6 \times N \times N}$ with a single model would introduce a high amount of redundancy across faces. A capacitance extraction typically requires millions of transitions, demanding a large number of neural network eval-

Algorithm 1: Random Dielectric Configuration

Input: Block count B , nesting probability p_{nest}

Output: A list of dielectrics L

```

1:  $L \leftarrow []$ 
2: for  $i = 1$  to  $B$  do
3:   Sample  $c_x, c_y, c_z \sim U(-2, 2)$ ;  $l_x, l_y, l_z \sim U(0, 4)$ ;
4:   Sample  $\kappa \sim \text{RandomDielectric}$ ;
5:   Append a block  $(c_x, c_y, c_z, l_x, l_y, l_z, \kappa)$  to  $L$ ;
6:   while  $U(0, 1) \leq p_{\text{nest}}$  do
7:      $l' \leftarrow \max(l_x, l_y, l_z)/10$ ;
8:      $l_x \leftarrow l_x + l'$ ;  $l_y \leftarrow l_y + l'$ ;  $l_z \leftarrow l_z + l'$ ;
9:     Sample  $\kappa \sim \text{RandomDielectric}$ ;
10:    Append a block  $(c_x, c_y, c_z, l_x, l_y, l_z, \kappa)$  to  $L$ ;
11:   end while
12: end for
13: Sample  $\kappa \sim \text{RandomDielectric}$ ;
14: Append a background block  $(0, 0, 0, \infty, \infty, \infty, \kappa)$  to  $L$ ;

```




Figure 3: Face selector network architecture ($N = 23$).

uations. This requirement makes model efficiency critical. Additionally, such a model would have to accurately predict both the shape of the Poisson kernel on each face as well as the overall probability across all faces.

A two-stage prediction system avoids this redundancy and simplifies the learning complexity by first predicting a categorical distribution $\mathbf{F} \in \mathbb{R}^6$ where $\mathbf{F}_i = \sum_{j,k} (\mathbf{p}_\alpha)_{i,j,k}$ across the six faces, and afterwards, predicting the conditional probabilities on the selected face $\frac{(\mathbf{p}_\alpha)_{i,:,:}}{\mathbf{F}_i} \in \mathbb{R}^{N \times N}$.

The categorical distribution \mathbf{F} requires modeling the interactions between the dielectrics within the transition cube. To capture these 3D spatial dependencies efficiently, we employ a 3D convolutional network referred to as the *face selector* \mathcal{F}_θ . It consists of four convolutional layers with a stride of 2 to progressively downsample the input dielectric tensor, as shown in Fig. 3. The extracted features are then processed by a feed-forward network. The model output is normalized via softmax. The model is trained to minimize the KL divergence loss:

$$\mathcal{L}_{\text{face-select}}(\theta) = D_{KL}[\mathbf{F} \parallel \text{softmax}(\mathcal{F}_\theta(\mathcal{X}))]. \quad (7)$$

For single-face Poisson kernel prediction, we treat the cube slices parallel to the target face as a stack of feature maps with N channels. The *face predictor* \mathcal{G}_θ is implemented using 2D depthwise separable convolutions (Howard et al. 2017), as shown in Fig. 4. The depthwise convolutions capture spatial patterns within each cube slice, while pointwise convolutions model inter-slice interactions across




Figure 4: Face predictor network architecture ($N = 23$).

the channel dimension. This separation significantly reduces computational cost and aligns with the fact that the Poisson kernel’s dependence on the dielectrics decays rapidly with distance from the surface.

To provide explicit spatial context and assist the learning of the low-frequency components of the Poisson kernel, we concatenate the input with a simple grid-based positional encoding (PE), adding two extra channels (x, y) (Liu et al. 2018). This results in an input tensor $\tilde{\mathcal{X}} \in \mathbb{R}^{(N+2) \times N \times N}$. Since Poisson kernel is inherently non-negative, the last layer of \mathcal{G}_θ is a ReLU activation. Finally, the output of \mathcal{G}_θ is normalized using L1 normalization to ensure a valid probability distribution. The loss function is defined as:

$$\mathcal{L}_{\text{face-predict}}(\theta) = D_{KL} \left[\frac{(\mathbf{p}_\alpha)_{i,:}}{\mathbf{F}_i^\nabla} \left\| \left\| \frac{\mathcal{G}_\theta(\tilde{\mathcal{X}})}{\|\mathcal{G}_\theta(\tilde{\mathcal{X}})\|_1} \right\| \right. \right]. \quad (8)$$

As discussed in Section 3.2, the first transition in each walk requires a special set of quantities: the *weight value* w_α , the *sign distribution* \mathbf{s}_α and the *gradient kernel* \mathbf{g}_α . In practice, different components (x, y , or z) of the gradient can be requested, depending on the orientation of the Gaussian surface at the first sample point.

We only learn the z -component and derive the other components via cube symmetries, i.e., applying corresponding rotations and reflections to the dielectric input. Similarly, we employ a two-stage process, with a gradient face selector and gradient kernel predictors.

The weight prediction constitutes a cumulative regression task similar to the face selector. We extend the output dimension of face selector \mathcal{F}_θ from 6 to 7 to simultaneously predict the face-wise distribution $\mathbf{F}_i^\nabla = \sum_{j,k} (\mathbf{g}_\alpha)_{i,j,k}$ and the weight magnitude $w_\alpha \in \mathbb{R}$. To accommodate this increased complexity, we scale up the network size appropriately (see Section 5.2). The training objective combines a KL divergence for face-wise probabilities with Mean Squared Error (MSE) for weight regression:

$$\mathcal{L}_{\text{grad-face-select}}(\theta) = D_{KL}[\mathbf{F}^\nabla \parallel \text{softmax}(\mathcal{F}_\theta(\mathcal{X})_{1:6})] + \lambda |w_\alpha - \mathcal{F}_\theta(\mathcal{X})_7|^2. \quad (9)$$

Given the structural similarity between the sign distribution \mathbf{s}_α and the gradient kernel \mathbf{g}_α , we propose to learn the signed gradient kernel $\mathbf{s}_\alpha \mathbf{g}_\alpha$ rather than modeling them separately. However, the gradient kernel is equivariant only to a subgroup of cube symmetries, necessitating at least two specialized face predictors to properly characterize the tangent (parallel to Gaussian surface) and normal (perpendicular to Gaussian surface) faces. The normal-face gradient

kernels exhibit challenging learning characteristics including sign ambiguities and dual-peaked distributions, motivating the use of deeper network architectures (see Section 5.2). Since the signed target includes both positive and negative values, we do not employ the ReLU activation in the face predictor \mathcal{G}_θ and use MSE as the training loss:

$$\mathcal{L}_{\text{grad-face-predict}}(\theta) = \left| \frac{(\mathbf{s}_\alpha \mathbf{g}_\alpha)_{i,:}}{\mathbf{F}_i^\nabla} - \frac{\mathcal{G}_\theta(\tilde{\mathcal{X}})}{\|\mathcal{G}_\theta(\tilde{\mathcal{X}})\|_1} \right|^2. \quad (10)$$

4.4 High-Throughput Inference Implementation

Our implementation targets high-throughput capacitance extraction using (1) an asynchronous producer-consumer architecture for continuous GPU utilization, (2) multi-instance model deployment for pipeline parallelism, and (3) custom CUDA kernels for efficient data processing.

Batched GPU Processing Architecture When a walker encounters a non-stratified transition domain, the system employs the lightweight models described in Section 4.3, following the two-stage pipeline shown in Figure 5. To maximize GPU utilization, transition domain sampling is performed in batches rather than individually. We implement a producer-consumer architecture where each thread manages a walker pool and submits transition tasks to a lock-free queue.

Multi-Instance Deployment To reduce latency, we deploy multiple model instances enabling simultaneous GPU processing and CPU-GPU data transfers. We instantiate one Poisson solver for every two walker threads and one gradient solver shared across all walkers. Data is processed as soon as available, preventing walker threads from stalling while waiting for GPU results.

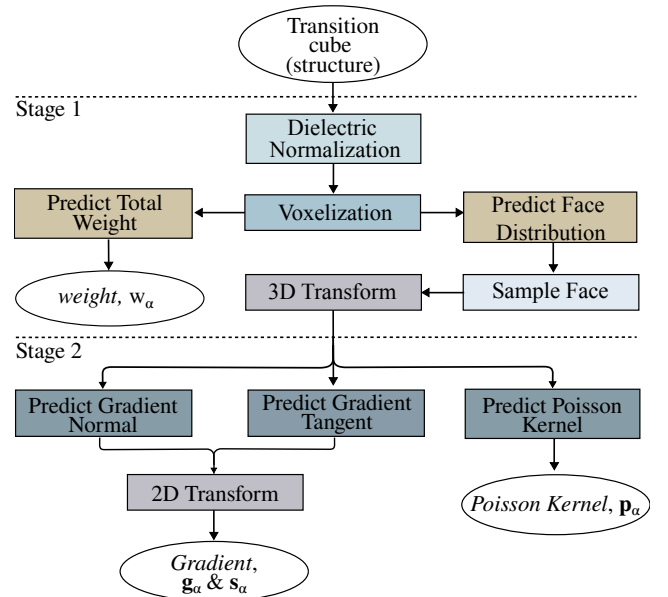


Figure 5: Transition quantities prediction pipeline.

Optimized Data Processing and Model Inference To maximize throughput, we minimize GPU memory transfers by transmitting compact structural descriptions rather than voxelized data, performing voxelization directly on GPU. Model inference uses TensorRT FP16 compilation and custom fused CUDA kernels for dielectric transformations.

5 Experimental Results

5.1 Experimental Setup

All experiments were conducted on a server equipped with an Intel Xeon Silver 4214 CPU @ 2.20GHz and an NVIDIA RTX 4090 GPU. The transition cubes were discretized with $N = 23$. The TensorRT compilation and custom CUDA kernel optimizations described in Section 4.4 are critical for achieving practical deployment speeds as shown in Figure 6.

5.2 Model Training

Datasets. We generated two comprehensive training datasets using the procedure described in Alg. 1 with $B = 5$ and $p_{\text{nest}} = 0.2$. The first dataset contains 100,000 samples for Poisson kernel prediction. The second dataset comprises 100,000 samples for gradient kernel prediction. Both datasets were split 90-10 for training and validation. The FDM solver took 1.7 hours to generate the dataset.

Training configuration. All models were trained using batch size of 16 for 200 epochs with the AdamW optimizer using $\beta_1 = 0.9$, $\beta_2 = 0.999$, and no weight decay. Gradient clipping was applied with a maximum norm of 1.0. The learning rate schedule employed cosine annealing from 1×10^{-3} to 5×10^{-6} with a warmup period during the first 20 epochs. For the dual-objective loss function in (9), we used $\lambda = 1$. Each full training run required 12.3 hours.

Model Implementation Details. The face solver architecture (see Fig. 4) comprises a positional encoding layer followed by a channel projection using 1×1 convolutions (from $23 + 2$ input channels to 16 channels), four depthwise separable layers with channel progression (16, 16, 8, 4, 2) and increasing dilation rates (1, 1, 2, 3), and a final 1×1 convolution head. The dilated convolutions enable the model to capture larger spatial contexts while maintaining computational efficiency and producing smoother output distributions. Each depthwise separable block consists of a 3×3 depthwise convolution, batch normalization, GELU activation, followed by a 1×1 pointwise convolution, batch normalization, and GELU activation. The tangent gradient solver variant follows the same structure, and the normal gradient solver includes wider channels (64, 64, 32, 32, 16, 16, 8, 4). The 3D face selector networks (see Fig. 3) use standard 3D convolutions with a stride of 2, batch normalization, GELU activations, followed by fully connected layers. We use channel progressions of (1, 4, 8, 16) for Poisson kernels and (1, 8, 16, 64) for gradient kernels.

5.3 Capacitance Extraction Results

Cases 1-6 involve simple parallel-wire structures across different metal layers and technologies. Cases 7-10 are FinFET structures with local interconnects, with cases 9-10 being

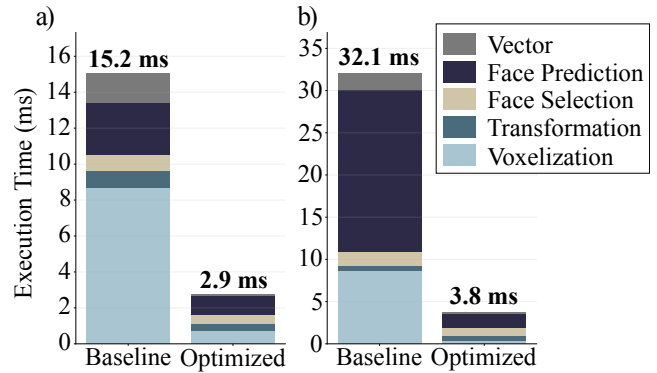


Figure 6: Performance analysis after introducing custom CUDA Kernels and TensorRT compilation, benchmarked with a batch size of 2048: (a) Poisson and (b) Gradient.

particularly challenging due to high-contrast conformal dielectrics. For all cases, the FRW process terminates when the estimated stochastic error falls below 1%. Ground-truth values are obtained from Raphael, which takes hours per case.

DeepRWCap was configured with 8 walker threads (512 walkers per thread), 4 threads for Poisson sampling, and 1 thread for gradient sampling. The CPU-based FRW-FDM, FRW-AGF (Huang and Yu 2024) and Microwalk (Huang and Yu 2025) used 16 walker threads. For a fair comparison, the GE-CNN approach (Visvardis, Liaskovitis, and Efsthathiou 2023) was adapted to the same GPU-accelerated architecture as DeepRWCap.

Table 2 compares method accuracy across test cases, with each method evaluated over 10 independent trials. The percentage of non-stratified domains requiring neural network inference varies significantly: Poisson kernel sampling in 11.2-55.8% of domains and gradient kernel sampling in 0.0-13.4% of domains. Runtime comparisons are shown in Figure 7. The FDM-guided FRW is omitted since it is over $10\times$ slower than DeepRWCap. Overall, cases 7-10 present shorter runtimes because the smaller conductor spacing accelerates FRW convergence. Statistical analysis using Wilcoxon signed-rank tests ($\alpha = 0.05$) shows DeepRWCap achieves a $1.23\times$ speedup over Microwalk ($p = 0.024$) while maintaining statistically equivalent accuracy. Although AGF achieved faster execution times, it exhibited poor reliability with high variance in error rates ($5.18\% \pm 7.81\%$ vs DeepRWCap’s $1.24\% \pm 0.53\%$).

Model	Params	FLOPs	Loss	Valid. Loss
Poisson Solver	1.40 K	0.84 M	KL	2.1×10^{-3}
Gradient Tangent	1.40 K	0.84 M	MSE	3.2×10^{-8}
Gradient Normal	5.96 K	3.48 M	MSE	4.5×10^{-8}
Poisson Selector	13.16 K	0.31 M	KL	7.6×10^{-4}
Gradient Selector	196.7 K	1.16 M	KL+MSE	2.8×10^{-3}

Table 1: Component Model Summary

Case	Node (nm)	FRW-FDM	GE-CNN	FRW-AGF	Microwalk	DeepRWCap (ours)			
		Error (%)	Error (%)	Error (%)	Error (%)	Error (%)	Tasks ^a (M)	Poisson ^b (%)	Grad. ^b (%)
1	16	0.4 ± 0.2	4.9 ± 0.2	0.8 ± 0.1	0.6 ± 0.2	1.2 ± 0.1	23.20	54.7	11.8
2	16	1.1 ± 0.3	<u>7.0 ± 0.4</u>	1.8 ± 0.3	1.7 ± 0.3	2.1 ± 0.3	5.79	27.4	4.4
3	16	1.0 ± 0.1	<u>5.6 ± 0.2</u>	2.2 ± 0.1	1.8 ± 0.2	1.2 ± 0.1	7.08	55.8	13.4
4	28	1.8 ± 0.4	<u>6.6 ± 0.3</u>	1.4 ± 0.4	0.5 ± 0.4	0.7 ± 0.3	41.07	27.7	2.6
5	55	1.0 ± 0.3	<u>3.9 ± 0.1</u>	0.7 ± 0.2	0.9 ± 0.3	2.2 ± 0.2	57.64	15.5	0.0
6	55	0.4 ± 0.2	<u>7.0 ± 0.2</u>	2.0 ± 0.3	0.4 ± 0.3	0.9 ± 0.3	41.67	26.3	0.0
7	16	2.2 ± 0.7	1.0 ± 0.3	1.7 ± 0.5	2.0 ± 0.6	1.2 ± 0.4	4.18	11.2	1.3
8	16	0.7 ± 0.6	0.9 ± 0.3	0.3 ± 0.3	0.6 ± 0.5	0.6 ± 0.5	2.33	10.9	1.5
9	12	1.5 ± 0.8	<u>22.9 ± 0.7</u>	<u>17.0 ± 1.0</u>	0.9 ± 0.5	1.1 ± 0.9	12.02	13.8	2.3
10	12	0.6 ± 0.6	<u>27.1 ± 0.8</u>	<u>23.9 ± 1.3</u>	0.6 ± 0.4	1.2 ± 0.9	6.23	14.3	3.1

Table 2: Performance Comparison of Random Walk Methods (with relative errors larger than 5% underlined).

^a Average number of transition cubes solved by DeepRWCap.

^b Average percentage of transition cubes classified as non-stratified, requiring neural inference for kernel sampling.




Figure 7: Method Runtime Comparison

Architecture	Params	FLOPs	L2 (%)	KL Div
MLP	34.4 M	34.4 M	7.83	0.0133
3D Conv	19.7 K	20.3 M	24.15	0.0298
2D Conv	4.28 K	2.31 M	13.80	0.0125
GE-CNN + GMM	0.43 M	0.65 M	26.63	0.0403
DS Conv	1.37 K	0.82 M	12.15	0.0083
DS Conv + Learn PE	2.46 K	0.84 M	4.11	0.0023
DS Conv + Grid PE	1.40 K	0.84 M	3.93	0.0021

Table 3: Model Architecture Ablation

5.4 Ablation Study

We conduct an ablation study to assess the impact of architectural design choices on accuracy and efficiency for single-face Poisson kernel prediction. We evaluate model variants under consistent training settings in Section 5.2. The MLP baseline includes three hidden layers of 2048 units each. We also evaluate with standard 2D convolutions. As a third baseline, we replace the depthwise separable convolution layers with 3D convolutional layers with stride-2 downsampling to reduce depth while preserving width and height.

From Table 3, depthwise separable convolutional models achieve superior accuracy with significantly fewer parameters and FLOPs compared to MLPs, 2D convolutions, and 3D convolutions. The 3D convolutional approach, while

conceptually appealing for processing volumetric dielectric data, proves computationally inefficient with substantially higher parameter count and FLOPs. This validates our design choice of decomposing the 3D problem into face selection and 2D face-specific kernel prediction. Moreover, incorporating positional encodings notably improves performance, with the fixed grid encoding yielding the best accuracy and lowest KL divergence.

6 Conclusion

We presented DeepRWCap, a neural-guided random walk solver that accelerates capacitance extraction for advanced IC designs. Our approach combines a two-stage CNN architecture with GPU-optimized inference to predict transition kernels in multi-dielectric domains. The compact depthwise separable networks with positional encoding achieve high accuracy while maintaining computational efficiency. Evaluated on 10 industrial test cases spanning 12-55 nm technologies, DeepRWCap demonstrates a mean relative error of 1.24% with significant speedups over state-of-the-art methods—achieving up to 49% acceleration on complex designs. The framework provides a practical solution for capacitance extraction in modern semiconductor technologies and serves as a basis for future extensions to other elliptic PDEs.

Acknowledgments

This work was partially supported by National Science and Technology Major Project (2021ZD0114703), Beijing Natural Science Foundation (Z220003), and the MIND project (MINDXZ202406). W. Yu is the corresponding author.

References

- Axler, S.; Bourdon, P.; and Wade, R. 2001. *Harmonic Function Theory*. Graduate Texts in Mathematics. Springer. ISBN 9780387952185.
- Cai, Y.; Liang, Y.; Luo, Z.; Xie, B.; and Li, X. 2024. PCT-Cap: Point Cloud Transformer for Accurate 3D Capacitance Extraction. In *Proceedings of the 2024 2nd International Symposium of Electronics Design Automation (ISED)*, 421–426. IEEE.
- Cheng, R.; Lyu, X.; Li, Y.; Ye, J.; Hao, J.; and Yan, J. 2022. The policy-gradient placement and generative routing neural networks for chip design. *Advances in Neural Information Processing Systems*, 35: 26350–26362.
- Gomes, W.; Koker, A.; Stover, P.; Ingerly, D.; Siers, S.; Venkataraman, S.; Pelto, C.; Shah, T.; Rao, A.; O’Mahony, F.; Karl, E.; Cheney, L.; Rajwani, I.; Jain, H.; Cortez, R.; Chandrasekhar, A.; Kanthi, B.; and Koduri, R. 2022. Ponte Vecchio: A Multi-Tile 3D Stacked Processor for Exascale Computing. In *IEEE International Solid-State Circuits Conference (ISSCC)*, 42–44. IEEE.
- Griffiths, D. J. 2023. *Introduction to Electrodynamics*. Cambridge University Press, 5 edition.
- Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *CoRR*, abs/1704.04861.
- Huang, J.; Yang, M.; and Yu, W. 2024. The Floating Random Walk Method With Symmetric Multiple-Shooting Walks for Capacitance Extraction. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 43(7): 2098–2111.
- Huang, J.; and Yu, W. 2024. Enhancing 3-D Random Walk Capacitance Solver with Analytic Surface Green’s Functions of Transition Cubes. In *Proc. DAC*.
- Huang, J.; and Yu, W. 2025. Efficient FRW Transitions via Stochastic Finite Differences for Handling Non-Stratified Dielectrics. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*
- Kakutani, S. 1944. Two-dimensional brownian motion and harmonic functions. *Proceedings of the Imperial Academy*, 20(10): 706–714.
- Lai, Y.; Lee, S.; Chen, G.; Poddar, S.; Hu, M.; Pan, D. Z.; and Luo, P. 2025. Analogcoder: Analog circuit design via training-free code generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 379–387.
- Li, Z.; Yang, G.; Deng, X.; De Sa, C.; Hariharan, B.; and Marschner, S. 2023. Neural caches for monte carlo partial differential equation solvers. In *SIGGRAPH Asia 2023 Conference Papers*, 1–10.
- Liebmann, L.; Smith, J.; Chanemougame, D.; and Gutwin, P. 2021. CFET design options, challenges, and opportunities for 3D integration. In *2021 IEEE International Electron Devices Meeting (IEDM)*, 3–1. IEEE.
- Liu, L.; Yang, F.; Shang, L.; and Zeng, X. 2024. GNN-Cap: Chip-Scale Interconnect Capacitance Extraction Using Graph Neural Network. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 43(4): 1206–1217.
- Liu, R.; Lehman, J.; Molino, P.; Petroski Such, F.; Frank, E.; Sergeev, A.; and Yosinski, J. 2018. An intriguing failing of convolutional neural networks and the coordconv solution. *Advances in neural information processing systems*, 31.
- Miller, B.; Sawhney, R.; Crane, K.; and Gkioulekas, I. 2024. Differential Walk on Spheres. *ACM Trans. Graph.*, 43(6).
- Mukesh, S.; and Zhang, J. 2022. A review of the gate-all-around nanosheet FET process opportunities. *Electronics*, 11(21): 3589.
- Muller, M. E. 1956. Some continuous Monte Carlo methods for the Dirichlet problem. *The Annals of Mathematical Statistics*, 569–589.
- Nam, H. C.; Berner, J.; and Anandkumar, A. 2024. Solving poisson equations using neural walk-on-spheres. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*.
- Øksendal, B. 2003. *Stochastic Differential Equations: An Introduction with Applications, Sixth Edition*. Springer.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378: 686–707.
- Sawhney, R.; and Crane, K. 2020. Monte Carlo geometry processing: A grid-free approach to PDE-based methods on volumetric domains. *ACM Trans. Graph.*, 39(4).
- Sawhney, R.; Seyb, D.; Jarosz, W.; and Crane, K. 2022. Grid-free Monte Carlo for PDEs with spatially varying coefficients. *ACM Trans. Graph.*, 41(4).
- Smolić, I.; and Klajn, B. 2021. CAPACITANCE MATRIX REVISITED. *Progress In Electromagnetics Research B*, 92: 1–18.
- Visvardis, M.; Liaskovitis, P.; and Efstathiou, E. 2023. Deep-learning-driven random walk method for capacitance extraction. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 42(8): 2643–2656.
- Worrall, D.; and Brostow, G. 2018. Cubenet: Equivariance to 3d rotation and translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 567–584.
- Wuu, J.; Agarwal, R.; Ciraula, M.; Dietz, C.; Johnson, B.; Johnson, D.; Schreiber, R.; Swaminathan, R.; Walker, W.; and Naffziger, S. 2022. 3D V-Cache: The Implementation of a Hybrid-Bonded 64MB Stacked Cache for a 7nm x86-64 CPU. In *IEEE International Solid-State Circuits Conference (ISSCC)*, volume 2022, 428–429. IEEE.

- Y. L. Le Coz and R. B. Iverson. 1992. A stochastic algorithm for high speed capacitance extraction in integrated circuits. *Solid-State Electron.*, 35(7): 1005–1012.
- Yang, D.; Li, H.; Yu, W.; Guo, Y.; and Liang, W. 2023. CNN-Cap: Effective Convolutional Neural Network Based Capacitance Models for Interconnect Capacitance Extraction. *ACM Transactions on Design Automation of Electronic Systems*, 28(4): 1–22.
- Yang, M.; and Yu, W. 2020. Floating random walk capacitance solver tackling conformal dielectric with on-the-fly sampling on eight-octant transition cubes. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 39(12): 4935–4943.
- Yu, W.; Song, M.; and Yang, M. 2021. Advancements and challenges on parasitic extraction for advanced process technologies. In *Proc. ASP-DAC*, 841–846.
- Yu, W.; Zhuang, H.; Zhang, C.; Hu, G.; and Liu, Z. 2013. RWCap: A floating random walk solver for 3-D capacitance extraction of very-large-scale integration interconnects. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 32(3): 353–366.
- Zhong, R.; Ye, J.; Tang, Z.; Kai, S.; Yuan, M.; Hao, J.; and Yan, J. 2024. Preroutggn for timing prediction with order preserving partition: Global circuit pre-training, local delay learning and attentional cell modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17087–17095.