

High-Dimensional Yield Estimation using Shrinkage Deep Features and Maximization of Integral Entropy Reduction

Shuo Yin*

School of Integrated Circuit Science
and Engineering, Beihang University
Beijing, China
18231082@buaa.edu.cn

Guohao Dai*

College of Mechatronics and Control
Engineering, Shenzhen University
Shenzhen, China
daiguohao2019@email.szu.edu.cn

Wei W. Xing^{†‡}

School of Integrated Circuit Science
and Engineering, Beihang University
Beijing, China
wxing@buaa.edu.cn

ABSTRACT

Despite the fast advances in high-sigma yield analysis with the help of machine learning techniques in the past decade, one of the main challenges, the curse of “dimensionality”, which is inevitable when dealing with modern large-scale circuits, remains unsolved. To resolve this challenge, we propose an absolute shrinkage deep kernel learning, ASDK, which automatically identifies the dominant process variation parameters in a nonlinear-correlated deep kernel and acts as a surrogate model to emulate the expensive SPICE simulation. To further improve the yield estimation efficiency, we propose a novel maximization of approximated entropy reduction for an efficient model update, which is also enhanced with parallel batch sampling for parallel computing, making it ready for practical deployment. Experiments on SRAM column circuits demonstrate the superiority of ASDK over the state-of-the-art (SOTA) approaches in terms of accuracy and efficiency with up to 11.1x speedup over SOTA methods.

KEYWORDS

Yield Analysis, Bayesian Optimization, Failure Probability

ACM Reference Format:

Shuo Yin, Guohao Dai, and Wei W. Xing. 2023. High-Dimensional Yield Estimation using Shrinkage Deep Features and Maximization of Integral Entropy Reduction. In *28th Asia and South Pacific Design Automation Conference (ASPDAC '23)*, January 16–19, 2023, Tokyo, Japan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3566097.3567907>

1 INTRODUCTION

As semiconductor fabrication technology improves by shrinking down its scale to nano-meter, the negative effect of the process variance, e.g., doping fluctuation, intra-die mismatches, and threshold voltage variation, arises and causes yield reduction. For circuits with cells replicated millions of times (e.g., SRAM), extremely small circuit failure probability (usually smaller than 10^{-6}) must be

considered to provide a robust design against fabrication process variations, which forms the yield analysis problem.

Monte Carlo (MC) analysis is generally considered the gold standard for yield analysis in industry and academia. To provide a reasonably accurate yield estimation, MC requires a large number (usually millions) of SPICE simulations, making it infeasible for modern yield problems, e.g., yield estimation for SRAM array with more than 500 independent process variation parameters. Taking a 32M SRAM with a 97% yield rate for an example, the yield of the bit cells needs to exceed 99.9999% [1]. Using MC, more than 1×10^6 MC samples are required to ensure accuracy. Therefore, academia and industry rely on other approximated yield estimation algorithms to reduce the overall time cost of the repeated simulations.

Instead of drawing samples randomly as in MC, Importance sampling (IS) based approaches draw samples according to a constructed distribution shifted to the likely-to-fail regions. For example, [2] shifts the sampling mean to the min-norm points of each failure region. [3] utilizes an adaptive resampling scheme to keep the sample mean updated. [4] combines variance reduction techniques with importance sampling and gate delay model. The convergence of probability estimation can be accelerated because the failure event is more likely to be drawn around the likely-to-fail regions. However, These IS-based methods can only find the nearest failure region, whereas the other failure regions are ignored, leading to low efficiency or even low accuracy when the number of samples is not sufficiently large. Surrogate-modeling-based methods construct a surrogate model to approximate the circuit simulators, based on which the yield is estimated. For example, [5] evaluates the circuit through a Gaussian process regression model, whereas [6] utilizes a polynomial chaos expansion model with low-rank tensor approximation to emulate the system performance function. To further improve a surrogate model’s efficiency, the training data is added sequentially based on the current estimations [6] rather than just relying on pre-sampling inputs.

It is also possible to combine the IS- and surrogate-based approaches. [5] uses an RBF neural network to fit the simulation results with an optimal-mean-shift-vector and proposes a heuristic algorithm to find candidates to conduct subsequent experiments. However, high-dimensional yield estimation is still a challenge for both industry and academia. For the surrogate-based method, the data-driven surrogate model (being Gaussian process or deep learning) does not scale well with the dimensionality of the input space due to the “curse of dimensionality” where the number of training data to cover the domain grows exponentially. Similarly, the IS-based method requires a large number of simulations to obtain

*Both authors contributed equally to this research.

[†] Corresponding author.

[‡] Also affiliated with Beihang Hangzhou Innovation Institute Yuhang, Hangzhou, China

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

ASPAC '23, January 16–19, 2023, Tokyo, Japan

© 2023 Association for Computing Machinery.

ACM ISBN 978-1-4503-9783-4/23/01...\$15.00

<https://doi.org/10.1145/3566097.3567907>

the likely-to-fail region, which suffers the same exponential growth with the increase of dimension.

To tackle the high-dimensional yield estimation challenge, We propose a novel deep kernel learning surrogate model with non-linear feature selection to capture the black box function between input and output; a parallel sampling scheme is also proposed to update the model efficiently. The novelty of our work includes: (1) We propose shrinkage deep features to enable the widely used GP surrogate for high-dimensional process variation inputs. Specifically, we introduce an absolute shrinkage in the reproducing kernel Hilbert Space (RKHS) via the Hilbert-Schmidt independence criterion (HSIC) to select the key features that dominate the input-output mapping, making a GP easier to train, robust against overfitting, and efficient to update with any acquisition functions. (2) We harness the rich model capacity of a deep kernel learning Gaussian process as our surrogate model to capture the complex black-box function of the process variation parameters and their circuit performance metric under a SPICE simulator. (3) We proposed a scalable parallel batch strategy to enable massive parallel model updates, which takes advantage of the high dimensionality and turns the “curse of dimensionality” into a “blessing of dimensionality”. (4) The empirical study shows that ASDK is up to 11.1x faster and more accurate than the SOTA methods.

2 BACKGROUND

2.1 Problem Definition: Rare Event Analysis

Define $\mathbf{x} = [x^{(1)}, x^{(2)}, \dots, x^{(d)}]^T \in X$ as the variational parameters, which explains the inevitable random variations of a manufacturing process when conducting a SPICE simulation with given design parameters, e.g., transistor widths and lengths, resistance values, capacitance values, and bias voltages and currents. Without loss of generality, \mathbf{x} is assumed independent Gaussian distributed after normalization, i.e.,

$$p(\mathbf{x}) = \prod_i^d \exp\left(-\frac{(x^{(i)})^2}{2}\right) / \sqrt{2\pi}. \quad (1)$$

For a particular \mathbf{x} (bare in mind that we cannot control the value of \mathbf{x}), the circuit performance metric z_k , e.g., amplifier gain and memory read/write time, can be considered as a function $z_k = f_k(\mathbf{x})$. When all K metrics are smaller/larger than some predefined threshold z^0 , e.g., $z_k \leq z_k^0$ for $k = 1, \dots, K$, the circuit with the corresponding parameters \mathbf{x} is considered a qualified design. Otherwise, it is a failure case. Let's use a compact notation $\mathbf{z} = \mathbf{f}(\mathbf{x})$ to denote the process, where \mathbf{f} includes all SPICE simulations and other necessary calculations. Putting this into a strict formulation, for a specific design circuit, the circuit failure probability (equivalently the yield rate) P_f is defined as

$$P_f \triangleq \int_X I(\mathbf{f}(\mathbf{x})) p(\mathbf{x}) d\mathbf{x}, \quad (2)$$

where $I : \mathbb{R}^k \rightarrow \{0, 1\}$ is the indicator function of whether a performance metric passes the predefined criteria. The integration is challenging as it does not admit a closed-form solution in general. In practice, we can generate N samples from $p(\mathbf{x})$ and approximate P_f by $P_f \approx \frac{1}{N} \sum_{i=1}^N I(\mathbf{f}(\mathbf{x}_i))$, which become exact when $N \rightarrow \infty$. Nevertheless, the computation of P_f is highly time-consuming as

the number of simulations N required to finish the integral is large particular for large d , where each simulation can take hours to finish.

2.2 Surrogate Model

To avoid frequent calls to the expensive SPICE simulators and possibly other calculations, we can use a data-driven surrogate model $\mathbf{g}(\mathbf{x})$ to approximate $\mathbf{f}(\mathbf{x})$ and use it to provide a quick estimation for any $\mathbf{x} \in X$. There are many choices for the surrogate model for the applications, such as RBF neural networks [5], polynomial chaos expansion [6], and Gaussian process [5]. The different methods have their own unique characteristics for specific scenarios. In general, with a surrogate model, we can approximate the yield using $\hat{P}_f = \frac{1}{N} \sum_{i=1}^N I(\mathbf{g}(\mathbf{x}_i))$, where \mathbf{x}_i can be obtained cheaply by sampling from the process variation distribution $p(\mathbf{x})$. The computation is fast because executing $\mathbf{g}(\mathbf{x}_i)$ is computationally cheap once trained. Note that $\mathbf{g}(\mathbf{x})$ is an approximation of $\mathbf{f}(\mathbf{x})$. Thus, \hat{P}_f is an approximation of P_f . The accuracy of \hat{P}_f depends on the $\mathbf{g}(\mathbf{x})$, which relies on the collected training dataset D . Certainly, we can use the design of experiment (DoE) to generate training inputs using Latin hypercube sampling (LHS) or Sobol sequence. However, this can be of low efficiency, particularly for the yield problem where only a few critical regions matters. Instead, we can use a sequential update scheme to update the surrogate model such that the surrogate always gets the best update in terms of reducing its error and uncertainty of P_f . This is in line with the theory of Bayesian optimization (BO), where to goal is to approach the global optimal by proposing a sequence of query points. Following BO, we call the function measuring the contribution of a candidate point \mathbf{x} an acquisition function (Acq). The sequential model updating can be formulated, $\mathbf{x}_* = \operatorname{argmax}_{\mathbf{x} \in X} \operatorname{Acq}(\mathbf{x}|D)$, where \mathbf{x}_* indicates the best candidates for the currently available data D , which is then updated using \mathbf{x}_* and its corresponding performance metric. Thanks to the acquisition function, we can find the failure region more efficiently without wasting our simulation in unnecessary regions, resulting in a speedup in estimating the yield.

3 RESEARCH METHODS

In this section, we present ASDK based on the deep kernel learning Gaussian process, feature selection, adaptive updating, and parallel acceleration for the yield estimation. The code is available on Github¹.

3.1 Deep kernel Learning Gaussian process

Gaussian process (GP) is a common choice as the surrogate model for design space exploration tasks due to its model accuracy and flexibility with uncertainty quantification. By giving the set of the process variation parameters and their circuit performance metric under the SPICE simulation, we aim to capture the black-box function of the process variational parameters and their circuit performance metric.

Suppose that we have a set of N yield observations $\{y_i = f(\mathbf{x}_i) + \epsilon_i\}_{i=1}^N$, where the noise ϵ is caused by the numerical error of a simulator and is assumed normally distributed. A GP model places a prior

¹<https://github.com/SawyerDust1228/HSIC-DKL-Yield-Estimation>

distribution over the function f as $f(\mathbf{x})|\theta \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'|\theta))$, where μ is the mean function, and the kernel function k is parameterized by θ . Usually, the mean function can be assumed zero, i.e., $\mu(\mathbf{x}) \equiv 0$, by centering the data. The kernel function can take many forms, the standard RBF kernel are as $k(\mathbf{x}, \mathbf{x}'|\theta) = \theta_0 \exp(-(\mathbf{x} - \mathbf{x}')^T \text{diag}(\theta_1, \dots, \theta_l)(\mathbf{x} - \mathbf{x}'))$. With this prior and available data $\mathbf{y} = (y_1, \dots, y_N)^T$, we can derive the model likelihood

$$\bar{L} = -\frac{1}{2}\mathbf{y}^T(\mathbf{K} + \sigma^2\mathbf{I})^{-1}\mathbf{y} - \frac{1}{2}\ln|\mathbf{K} + \sigma^2\mathbf{I}| - \frac{N}{2}\log(2\pi), \quad (3)$$

where $\mathbf{K} = [K_{ij}]$ represent the covariance matrix, in which $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, $i, j = 1, \dots, N$, and σ^2 denotes the variance of the noise ϵ . The hyperparameters θ are normally obtained from point estimates by maximum likelihood estimate (MLE) w.r.t. θ .

Recently, deep neural networks have achieved great success in many areas because of their remarkable capacity for feature extraction. [7] combines the non-parametric flexibility of a kernel function with the powerful model capacity of the deep neural networks, which significantly improves the performance of a general GP. Despite its success, the underlying ideas are rather simple. It essentially redefines the kernel as

$$\hat{k}(\mathbf{x}, \mathbf{x}') = k(\phi(\mathbf{x}, \mathbf{w}), \phi(\mathbf{x}', \mathbf{w})), \quad (4)$$

where $\phi(\mathbf{x}, \mathbf{w})$ is a deep neural network (for instance, a multi-layer perception (MLP) with multiple hidden layers) parameterized by weights \mathbf{w} and $k(\cdot, \cdot)$ is any valid kernel function, e.g., RBF kernel.

3.2 Shrinkage Deep Feature Selection

In general, the process variation parameters quantify the process corner and other factors during the fabrication of a circuit, e.g., threshold voltage, channel length modulation effect, and bulk effect. Usually, the same type of variational variables is applied to each transistor (and/or other crucial elements in the circuit). Due to the large number of transistors a practical circuit can have, we will end up with a large number of variational parameters, e.g., 1000, causing the ‘‘curse of dimensionality’’ issue, which makes both the integration and surrogate fitting extremely challenging.

Fortunately, previous research [8] reveals that not all variational parameters are equally important. In fact, circuit performance is dominated by several critical transistors, whereas the other transistors have little influence, especially in the case where the circuit has a special symmetric design to alleviate the process variation. This finding makes ‘‘dimension reduction’’ possible to reduce the input dimension such that only the key parameters are preserved. Note that this is more of a ‘‘feature selection’’ rather than dimension reduction because the inputs are fully independent, indicating that no dimension reduction techniques, e.g., PCA and KPCA, can achieve any success.

Let $\{x_i, y_i\}_{i=1}^N$ denotes the collection of all variation parameters and corresponding performance metric $y \in \mathbb{R}^D$. The surrogate model aims to approximate $f(\mathbf{x})$ with a small N and a large D , which can easily lead to overfitting. As mentioned above, only a few variation parameters matter the most. We consider the classic feature selection method, least absolute shrinkage selection operator (LASSO) [9], which chooses the key feature that best explains the input-output relationship. Essentially, LASSO casts an

optimization problem,

$$\operatorname{argmin}_{\alpha} \frac{1}{2}\|\mathbf{y} - \mathbf{X}\alpha\|^2 + \lambda\|\alpha\|^1, \quad (5)$$

where $\alpha \in \mathbb{R}^D$ is a weight vector determining the contribution of each dimension of \mathbf{x} , λ is the penalty factor; $\|\cdot\|^2$ denotes the L2 norm and $\|\cdot\|^1$ the L1 norm. This formulation is effective because its gradient w.r.t $\alpha^{(i)}$ (which is the i -th element of α) is a constant, which will push the value of $\alpha^{(i)}$ towards zero unless it contributes significantly to reducing the data fitting loss of the first term.

Despite its elegance and effectiveness, we can immediately tell the limitation of LASSO is that it is based on a linear model, whereas the yield analysis problem we are facing is usually nonlinear. An ordinary LASSO will ignore the nonlinear connections and only focus on the linear ones, leading to a poor feature selection. Since a GP essentially relies on the reproducing kernel Hilbert space (RKHS) to build the connection between the input and output, we can naturally extend the original LASSO for the RKHS, which leads us to the Hilbert-Schmidt independence criterion (HSIC) proposed by [10]. Specifically, we require the output correlation measured by the kernel function $k(\cdot, \cdot)$ can be captured via the sum (with weight factor α) of the correlation matrix of each input dimension. The weight factor α is then contrasted using the LASSO trick. For a more rigorous derivation, the readers are referred to [10]. In practice, we also apply double centering for the kernel matrix to ensure stability and solve the optimization problem:

$$\operatorname{argmin}_{\alpha} \frac{1}{2}\|\tilde{\mathbf{L}} - \sum_{d=1}^D \mathbf{K}^{(d)}\alpha^{(d)}\|^2 + \lambda\|\alpha\|^1 \quad (6)$$

where $\tilde{\mathbf{L}} = \mathbf{H}\mathbf{K}_y\mathbf{H}^T$, with $\mathbf{H} = \mathbf{I} - \frac{1}{N}\mathbf{1}\mathbf{1}^T$ being the centering matrix; $[\mathbf{K}_y]_{ij} = k(y_i, y_j)$ and $[\mathbf{K}^{(d)}]_{ij} = k(\mathbf{x}_i^{(d)}, \mathbf{x}_j^{(d)})$ are kernel matrix given the target values and d -dimension of \mathbf{x} ; $k(\cdot, \cdot)$ is the kernel function with default hyper-parameters. The key result of this algorithm is the weight vector α , which indicates the importance of each input dimension that dominates the input-output mapping in the RKHS, which is encoded in our GP.

3.3 Maximum Integral Entropy Reduction

In Bayesian Optimization, the acquisition usually uses expected improvement (EI), predictive improvement (PI), and upper confidence bound (UCB) to propose a candidate for the next iteration. However, they are designed for optimization and do not generalize to yield estimation directly.

We notice that the circuit yield is usually very high (equivalently, the failure rate is very low), which means that the passing threshold z^0 is an extreme value compared to most simulation performances, and only a few samples will eventually fail the indication function. This hints that we can reduce the computing cost of integral by avoiding observation in the region that the simulation performance z will ‘‘absolutely’’ pass the criterion and try to locate the boundary of the failure region. This search scheme is in line with the boundary search method.

Considering a GP $f(\mathbf{x}) = [f_0(\mathbf{x}), f_1(\mathbf{x}) \dots f_K(\mathbf{x})]^T$, the posterior of each $f_i(\mathbf{x}) \sim \mathcal{GP}(m_i(\mathbf{x}), k_i(\mathbf{x}, \mathbf{x}'))$ is also a gaussian distribution with mean $m_i(\mathbf{x})$ and variance $v_i(\mathbf{x})$. The predictive posterior based on the pass threshold z^0 is a Bernoulli distribution with likelihood

$l(\mathbf{x}) \triangleq p(\tilde{I}(\mathbf{x}) = 1)$, where $\tilde{I}(\mathbf{x}) = I(\mathbf{f}(\mathbf{x}))$ indicates the predictive performance,

$$l(\mathbf{x}) = \prod_{k=1}^K p(\tilde{f}_k(\mathbf{x}) \geq z_k) = \prod_{k=1}^K \Phi\left(\frac{\mu_k(\mathbf{x}) - z_k^0}{v_k(\mathbf{x})}\right), \quad (7)$$

where $\Phi(\cdot)$ is the cumulative density function (CDF) of a normal distribution. According to the Poisson binomial distribution, which is the sum of independent yes/no experiments. We derive the approximated integral \tilde{g} in a similar manner,

$$\begin{aligned} \mathbb{E}[\hat{P}_f] &= \int_{\mathcal{X}} l(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}, \\ \text{Var}[\hat{P}_f] &= \int_{\mathcal{X}} l(\mathbf{x}) (1 - l(\mathbf{x})) p(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (8)$$

Because $l(\mathbf{x})$ is tractable, and $p(\mathbf{x})$ is simply a diagonal Gaussian, we can compute the integral efficiently using numerical approximations or quasi-MC. Achieving an accurate yield estimation \hat{P}_f is equivalent to reducing its variance $\text{Var}[\hat{P}_f]$, which, however, ignores the high-order moment and can lead to inferior results. Instead of reducing the variance, we introduce a probability information entropy for the yield posterior of $\tilde{I}(\mathbf{x})$, which is the entropy of a Bernoulli distribution,

$$H(\mathbf{x}) = -l(\mathbf{x}) \log(l(\mathbf{x})) - (1 - l(\mathbf{x})) \log(1 - l(\mathbf{x})). \quad (9)$$

We then define the total integral entropy as

$$IH = \int_{\mathcal{X}} H(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}, \quad (10)$$

which indicates the uncertainty of $g(\mathbf{x})$ based on the surrogate model with current observations D . To reduce the uncertainty of P_f , we can propose a candidate base on maximizing the expected integral entropy reduction,

$$\begin{aligned} \mathbf{x}^* &= \underset{\mathbf{x} \in \mathcal{X}}{\text{argmax}} (IH(\mathcal{D}) - IH(\mathcal{D} \cup \mathbf{x})) \\ &= \underset{\mathbf{x} \in \mathcal{X}}{\text{argmin}} IH(\mathcal{D} \cup \mathbf{x}). \end{aligned} \quad (11)$$

Thus, to get the optimal \mathbf{x}^* , we first draw multiple samples from the predictive posterior $f_k(\mathbf{x}|\mathcal{D})$ for a possible observation performance metric $\mathbf{z}^{(k)}$. We then combine those samples with our data collection \mathcal{D} and update the posterior as $f_k(\mathbf{x}|\mathcal{D} \cup \{\mathbf{x}, \mathbf{z}^{(k)}\})$, based on which, we can compute the integral entropy $IH(\mathcal{D} \cup \mathbf{x})$.

3.4 Parallel Batch Acquisition

In practical yield applications, it is important to allow a parallel acquisition of multiple candidates to unleash the power of the modern cluster center. In a recent work [6], the authors use a mixed Gaussian distribution to implement a parallel updating scheme. More specifically, they generate a discrete Gaussian sampling distribution around each observed data and use it to sample new candidates. However, these new candidates may fall in the same region with a high probability in the input domain, leading to an inefficient sampling strategy. In this section, we generalize ASDK for parallel computing. More specifically, we would like to propose multiple candidates at each iteration. Formally, we aim to solve the optimization,

$$\mathbf{X}^* = \underset{\mathbf{X} \in \mathcal{X}}{\text{argmin}} IH(\mathcal{D} \cup \mathbf{x}_1 \cup \dots \cup \mathbf{x}_Q), \quad (12)$$

where \mathbf{X}^* indicates the collection of Q ideal query points. As discussed previously, directly solving this equation is challenging and computationally expensive. To approximately solve this optimization in a batch fashion, we take advantage of the ‘‘curse of dimensionality’’ and turn it into a blessing. More specifically, we discover that if we start the ASDK at multiple far-away initial locations, the final query points do not converge to the same locations due to the complex geometry in the high-dimensional space. Inspired by the Q -batch initialization [11], we convert finding multiple query points into finding multiple far-away initial points for optimization.

Assume that we aim to generate Q query points. We first pre-sample T points in the domain, where we ensure $T \gg Q$. We then compute these T points’ acquisition score \mathbf{s} , $\{s_i = IH(\mathcal{D} \cup \mathbf{x}_i)\}_{i=1}^T$. Let s_m be the max value in these T scores. Define γ as a fraction of the maximum observed value under which we will ignore. This coefficient helps us filter out the scores that are lower than $\gamma * s_m$. In cases where the number of satisfying candidates is smaller than the query number Q , we use the following equation to relax the threshold $\gamma = (1 - \beta)\gamma_0$, where β is another hyperparameter to relax the filter and γ_0 is the ratio that does not satisfy the need. We can repeat this process until we get more than O points ($O > Q$) over the threshold. This idea is similar to the two-stage estimation in [12]. After we get O samples $\{s_i\}_{i=1}^O$ that are larger than the threshold, we use the following equation to get the sample weights:

$$\omega_i = \exp(\eta_1 \frac{s_i}{s_m}), \quad (13)$$

where η_1 is a scaling factor. Based on the weight ω , we sample Q initial points sequentially from the K points set.

We then conduct maximum entropy reduction in Eq. (11) in parallel with the Q initial points with gradient descent, e.g., SGD or Adam. A summary of our parallel batch query scheme is presented in Algorithm 1.

Algorithm 1 ASDK Parallel Batch query Algorithm

- Input:** number of initial set T , number of candidate Q , number of filtered samples O , fraction coefficient γ , relax coefficient β , weight sample coefficient η_1 .
- 1: Randomly generate T points in the domain using the Sobel sequence and compute the scores for these points.
 - 2: Assign the threshold α_{max} and get O candidate points.
 - 3: Sequentially sample Q points with weight $\omega_i = \exp(\eta_1 \frac{s_i}{s_m})$
 - 4: Conduct maximum entropy reduction in Eq. (11) in parallel with the Q initial points to get Q optimal solutions \mathbf{X}^*
 - 5: **return** Best candidates \mathbf{X}^*
-

4 EXPERIMENT RESULTS

In this section, we assess ASDK with the SOTA yield estimation algorithms on commonly used high-dimensional benchmark circuits. We compare ASDK to the SOTA yield estimation algorithms, including: (1) LRTA [6], a high-dimensional yield estimation algorithm using the low-rank tensor approximate polynomial chaos expansion model as the surrogate model with a KDE-based adaptive sampling strategy to update the model, (2) HDBO [13], which uses a random embedding feature reduction method and a Bayesian optimization to find the failure event, (3) HSCS [2], which applies a




Figure 1: The schematic of 6T-SRAM cell (left) and SRAM array column (right)

clustering algorithm to identify multiple failure regions and uses min-norm-points to resample the failure region. Bayesian-based approaches like HDBO and ASDK are implemented with parallel computing. All experiments are performed on a Linux system with AMD 5950x, GTX 3080, and 32GB RAM.

To determine when to stop the yield estimation process, we follow the widely used Figure of Merit (FOM) ρ in the yield estimation literature [2, 3, 6] as the stopping criteria. FOM is defined by $\rho = \sigma_{P_f} / P_f$, where P_f denotes the mean failure probability estimation and σ_{P_f} the standard deviation of P_f . Following the literature, we set the threshold $\rho_0 = 0.1$, i.e., stopping the yield estimation process when $\rho < 0.1$. This is equivalent to the stopping criteria of convergence with 90% confidence interval[2].

Without loss of generality, we test one circuit metric for each experiment, i.e., $z^0 \in \mathbb{R}$. We can implement high dimensional circuit metric by changing our GP model to a multi-task GP [14] or just fitting each metric independently. For the experimental purpose, we set the threshold to fix the yield failure rate to be approximately 10^{-4} to reduce overall computation and to emphasize the search of multiple failure regions in the process variational space. The hyperparameters we mention in this section 3.4 is set as following, $\alpha = 0.3$, $\beta = 0.9$, $T = 100 \times N$, $\eta_1 = 0.5$. As for the deep kernel network, we use a three-layer MLP when the dimensionality is smaller than 128 and a four-layer MLP for other cases. The proposed method can be applied to general scenarios by changing the structure of the neural networks.

4.1 18 Dimensional 6T SRAM Bit Cell

The 6T SRAM bit cell is a basic storage unit of SRAM circuits, which consists of six transistors shown in Fig. 1. In the circuit design, M1, M3, M5, and M6 are NMOS transistors, whereas M2 and M4 are PMOS transistors, BL is the bit line, and WL is the word line. The state of each bit in the SRAM is stored in two cross-connected inverters composed of M1, M2, M3, and M4. M5 and M6 NMOS are control switches used to control data transmission from the storage unit to the in-place line.

We consider the delay of read/write as the circuit output metric for this experiment, which is a commonly used setup for yield estimation research [3, 6, 15]. In the setting of this experiment, our deep kernel learning uses a three layers MLP with 200, 100, and 10 hidden neurons and ReLu as the activation function between each layer. We use a Matern plus a linear kernel for the GP to keep flexibility and efficiency. Fig. 2 shows the detailed estimation process whereas Table 1 concludes the experimental results. It can be seen clearly from the table that ASDK outperforms the competitors with a large margin in terms of estimation accuracy and the number of




Figure 2: P_f and FOM on 18-dimensional 6T SRAM

simulations required to achieve convergence, rendering a 196.30x speedup compared to standard MC and an up to 6x speedup to the competitors. The proposed method is the most accurate with only 4.14% relative error w.r.t. the golden truth MC value among the baseline methods. The MC method reaches the stopping criteria after 265000 simulations. HSCS, HDBO, and LRAT converge with 8100, 3500, and 2200 simulations, respectively. The failure rate of the IS-based HSCS remains small in the beginning, and a increase in the failure rate will emerge when HSCS discovers correct failure regions. The poor efficiency of HSCS is not a surprise, because an IS-based approach needs a massive number of pre-sampling data to cover the failure regions. HDBO, LRAT, and the proposed ASDK exhibit relatively fast convergence with a 75.71x, a 120.45x and a 196.30x speedup because they use surrogates. Nevertheless, ASDK converge to the most accurate fail rate using merely 1350 simulations. The evolution trend of the yield for surrogate methods is determined by the initial sampling technique and the regression model. Due to the different initial fitting techniques to the variation parameter space, the initial failure rates of different surrogate approaches vary. With the surrogates being updated, the fail rates of LRAT and ASDK grow larger gradually, whereas that of HDBO grows smaller. Nevertheless, they all converge to the ground truth with enough simulation runs.

The time consumed in model training (the time spent in simulation run is not included) during the whole yield estimation process of each approach is shown in Table 1. MC needs no training. IS-based HSCS spends only 5.28s on training. Surrogate-based methods like HDBO, LRAT, and ASDK require more time in model training. Among them, HDBO and ASDK obviously need more time to train because their Bayesian optimization is time-consuming. The proposed ASDK require the most time to train in this case. However, the train time of the model can be ignored when compared with the time of the simulations.

4.2 569 Dimensional SRAM Column


Fig. 1 shows the simplified circuit diagram of the 6T-SRAM bit cell array. Similarly, we use the delay of read/write as the metric. We increase the dimensionality with more cells, which leads to 569 process variation parameters. When the dimension get large, we need to increase our deep network's capacity to deal with the more complex data structure. To this end, we use an MLP with 1000,

Table 1: Final P_f estimation on 18-dimensional 6T SRAM

	MC	HSCS	HDBO	LRTA	Proposed
Failure prob.	4.83e-4	5.15e-4	6.25e-4	6.40e-4	4.60e-4
Relative error	Golden	6.62%	29.40%	19.46%	4.14%
# of Sim.	265000	8100	3500	2200	1350
Sim. speedup	1x	32.72x	75.71x	120.45x	196.30x
Training time	N/A	5.28s	401.62s	53.50s	1537.73s


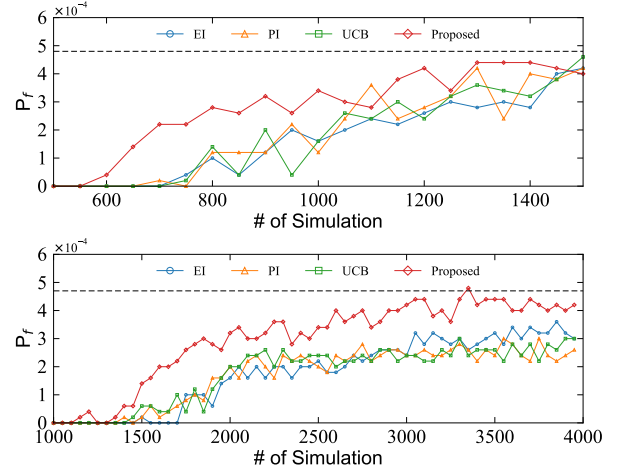
Table 2: Final P_f on 569-dimensional SRAM column

	MC	HSCS	HDBO	LRTA	Proposed
Failure prob.	4.70e-4	5.82e-4	3.87e-4	5.60e-4	4.39e-4
Relative error	Golden	23.83%	17.66%	19.14%	6.60%
# of Sim.	928500	44400	6100	5400	4000
Sim. speedup	1x	20.91x	152.21x	171.94x	232.13x
Training time	N/A	112.53s	1001.73s	12403.21s	5546.56s


Figure 3: P_f and FOM on 569-dimensional SRAM column

500, 200, and 20 hidden units with the hyperparameters remaining unchanged. The number of selected features is set at 120.

The same experimental results are shown in Fig. 3 and Table 2. The fail rate of HSCS remains very low until it finds failure regions near the origin of the coordinates of the variational parameter space. Therefore, the first half of its P_f curve is unable to be plotted in the figures. Surrogate-based methods like HDBO, LRTA, and ASDK converge to the ground truth gradually. We can see that ASDK again achieves the lowest relative error among all methods with the minimum number of simulation runs, rendering its superiority over the SOTA methods. Particularly, compared with the second best method, LRTA, ASDK is 2.9x more accurate and 1.3x faster. The improvement over HSCS is about 11.1x faster with a 3.6x improvement in accuracy. As for the training time, LRTA requires 12403s, which makes it the lowest efficient method in this case due to the exponential growth of complexity w.r.t. dimensionality. Notice that ASDK also requires quite a significant training time. However, we believe that this computational cost is weightless than the simulation runs.


Figure 4: P_f estimation with different batch size

Figure 5: Acquisition function experiment

4.3 Ablation Study

Parallel Batch Update Convergence Validation. To assess the proposed parallel batch update method, we conducted the parallel experiment on both previous experiments with different batch sizes, i.e., {20,50,100} and show the yield in Fig. 4. We can see that a parallel run with 100 simultaneous candidates is almost as good as with 20, indicating the scalability of the proposed batch method.

Maximum Integral Entropy Infill Validation. To demonstrate the superiority of the proposed surrogate updating method of ASDK over the commonly used EI, PI, and UCB acquisition functions, we run ASDK with different methods and compare the estimated P_f in Fig. 5 for both previous experiments. Despite that all methods converge to the ground truth, the proposed Maximum Integral Entropy Infill outperforms the competitors in terms of the convergence rate w.r.t the number of simulations in both cases, which is essential in yield analysis to save the expensive simulation cost.




Figure 6: Feature reduction experiment

Feature selection Validation. To demonstrate that the feature selection by the HSIC-Lasso algorithm is reliable and efficient, we compare ASDK with some traditional feature selection and dimension reduction algorithms, including Factor Analysis (FA), Principal Component Analysis (PCA), Mutual Information (MI) [8], and Random Embedding (RE) [13]. We first randomly generate 1000 training data through Latin hypercube sampling (LHS). We run the feature selection/dimension reduction algorithms to reduce the dimensionality to a specified number and assess the model accuracy using RMSE using left-out testing points. The experiments are repeated five times with randomly shuffled training and testing data, and the average RMSE against the number of preserved dimensionality is shown in Fig. 6. We can see very clearly that the implemented HSIC-Lasso is significantly more efficient in preserving the dominant information, leading to a much lower RMSE until the dimension reaches 500. This also partly explains the previous superior performance of ASDK.

5 CONCLUSION

In this paper, we propose ASDK, a shrinkage deep kernel learning to efficiently model high-dimensional input with an effective parallel batch updating scheme, to efficiently tackle the yield estimation of high-dimensional variational variables. Compared to the SOTA methods, ASDK shows nearly 11.1x speed up with consistently accurate results. Limitation of ASDK includes (1) solving entropy-based optimization relies on high-performance GPU and is time-consuming, and (2) the feature selection can ignore critical variables and lead to model bias for the yield estimation.

6 ACKNOWLEDGMENTS

This work is supported by the Fundamental Research Funds for the Central Universities. The experiment is supported by Primarius Technologies Co., Ltd.

REFERENCES

- [1] Lorenzo Ciampolini, Jean-Christophe Lafont, Faress Tissafi Drissi, Jean-Paul Morin, David Turgis, Xavier Jonsson, Cyril Descleves, and Joseph Nguyen. Efficient yield estimation through generalized importance sampling with application to nbl-assisted sram bitcells. In *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–8, 2016.
- [2] Wei Wu, Srinivas Bodapati, and Lei He. Hyperspherical clustering and sampling for rare event analysis with multiple failure region coverage. In *Proceedings of the 2016 on International Symposium on Physical Design*, pages 153–160, 2016.
- [3] Xiao Shi, Fengyuan Liu, Jun Yang, and Lei He. A fast and robust failure analysis of memory circuits using adaptive importance sampling method. In *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2018.
- [4] Alp Arslan Bayrakci, Alper Demir, and Serdar Tasiran. Fast monte carlo estimation of timing yield with importance sampling and transistor-level circuit simulation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 29(9):1328–1341, 2010.
- [5] Jian Yao, Zuoqiang Ye, and Yan Wang. An efficient sram yield analysis and optimization method with adaptive online surrogate modeling. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 23(7):1245–1253, 2015.
- [6] Xiao Shi, Hao Yan, Qiancun Huang, Jiajia Zhang, Longxing Shi, and Lei He. Meta-model based high-dimensional yield analysis using low-rank tensor approximation. In *Proceedings of the 56th Annual Design Automation Conference 2019*, pages 1–6, 2019.
- [7] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial intelligence and statistics*, pages 370–378. PMLR, 2016.
- [8] Jinyuan Zhai, Changhao Yan, Sheng-Guo Wang, and Dian Zhou. An efficient bayesian yield estimation method for high dimensional and high sigma sram circuits. In *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*, pages 1–6, 2018.
- [9] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- [10] Makoto Yamada, Wittawat Jitkritum, Leonid Sigal, Eric P. Xing, and Masashi Sugiyama. High-dimensional feature selection by feature-wise kernelized lasso. *Neural Computation*, 26(1):185–207, 2014.
- [11] Rommel G. Regis and Christine A. Shoemaker. Combining radial basis function surrogates and dynamic coordinate search in high-dimensional expensive black-box optimization. *Engineering Optimization*, 45(5):529–555, 2013.
- [12] Mengshuo Wang, Wenlong Lv, Fan Yang, Changhao Yan, Wei Cai, Dian Zhou, and Xuan Zeng. Efficient yield optimization for analog and sram circuits via gaussian process regression and adaptive yield estimation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 37(10):1929–1942, 2018.
- [13] Hanbin Hu, Peng Li, and Jianhua Z Huang. Enabling high-dimensional bayesian optimization for efficient failure detection of analog and mixed-signal circuits. In *2019 56th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2019.
- [14] Edwin V Bonilla, Kian Chai, and Christopher Williams. Multi-task gaussian process prediction. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [15] Xiao Shi, Hao Yan, Chuwen Li, Jianli Chen, Longxing Shi, and Lei He. A non-gaussian adaptive importance sampling method for high-dimensional and multi-failure-region yield analysis. In *2020 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, pages 1–8, 2020.