DNN-Opt: An RL Inspired Optimization for Analog Circuit Sizing using Deep Neural Networks

Ahmet F. Budak[1*], Prateek Bhansali[2], Bo Liu[3], Nan Sun[1], David Z. Pan[1] and Chandramouli V. Kashyap[2†]
[1]ECE Department, The University of Texas at Austin [2]Intel Corp. [3]James Watt School of Eng., University of Glasgow
*ahmetfarukbudak@utexas.edu, †chandramouli.v.kashyap@intel.com

***Abstract*—Analog circuit sizing takes a significant amount of manual effort in a typical design cycle. With rapidly developing technology and tight schedules, bringing automated solutions for sizing has attracted great attention. This paper presents DNN-Opt, a Reinforcement Learning (RL) inspired Deep Neural Network (DNN) based black-box optimization framework for analog circuit sizing. The key contributions of this paper are a novel sample-efficient two-stage deep learning optimization framework leveraging RL actor-critic algorithms, and a recipe to extend it on large industrial circuits using critical device identification. Our method shows 5–30x sample efficiency compared to other black-box optimization methods both on small building blocks and on large industrial circuits with better performance metrics. To the best of our knowledge, this is the first application of DNN-based circuit sizing on industrial scale circuits.**

***Index Terms*—Analog Circuit Sizing Automation, Blackbox Optimization, Reinforcement Learning, Deep Neural Network**

I. Introduction

Analog Integrated Circuit (IC) design is a complex process involving multiple steps. Billions of nanoscale transistor devices are fabricated on a silicon die and connected via intricate metal layers during those steps. The final product is an IC, which powers much of our life today. An essential aspect of IC design is analog design, which continues to suffer from long design cycles and high design complexity due to lack of automation in analog Electronic Design Automation (EDA) tools compared to digital flows. In particular, "circuit sizing" tends to consume a significant portion of analog designers' time. In order to tackle this labor-intensive nature and reduce time-to-market requirements, analog circuit sizing automation has attracted high interest in recent years.

Prior work on analog circuit sizing automation can be divided into two categories: knowledge-based and optimization-based methods. In the knowledge-based approach, design experts transcribe their domain knowledge into algorithms and equations [1], [2]. However, such methods create dependency on expert human-designers, circuit topology, and technology nodes. Thus, these methods are highly time-consuming and not scalable.

Optimization-based methods are further categorized into two classes: equation-based and simulation-based methods. Equation-based methods try to express circuit performance via posynomial equations or regression models using simulation data. Then the equation-based optimization methods such as Geometric Programming [3], [4] or Semidefinite Programming (SDP) relaxations [5] are applied to convex or non-convex formulated problems to find an optimal solution. Although those methods are generally fast, developing accurate expressions for circuit performances is not easy and deviates largely from the actual values. On the other hand, simulation-based methods employ black-box or learning-based optimization techniques to explore design space. These methods make guided exploration in the search space and target a global minimum using the real evaluations from circuit simulators.

Traditionally, there have existed various model-free optimization methods such as particle swarm optimization (PSO) [6] and advanced differential evolution [7]. Although these methods have good convergence behavior, they are known to be sample-inefficient (i.e., SPICE simulation intensive). Recently surrogate model-based and learning-based methods are becoming increasingly popular due to their efficiency in exploring solution space. In surrogate model-based methods, Gaussian Process Regression (GPR) [8] is generally used for design space modeling, and the next design point is determined through model predictions. For example, GASPAD method is introduced into Radio Frequency (RF) IC synthesis where GPR predictions guide evolutionary search [9]. WEIBO method proposed a GPR based Bayesian Optimization [10] algorithm where a blended version of weighted Expected Improvement (wEI) and the probability of feasibility is selected as acquisition function to handle constrained nature of analog sizing [11]. The main drawback of Bayesian Optimization methods is scalability as GP modeling has cubic complexity in the number of samples, $\mathcal{O}(N^3)$.

Recently, reinforcement learning algorithms are applied in the area as learning-based methods. GCN-RL [12] leverages Graph Neural Networks (GNN) and proposes a transferable framework. Despite reporting superior results over various methods and human-designer, a) it requires thousands of simulations for convergence (without transfer learning) and b) it suffers from engineering effort to determine observation vector, architecture selection, and reward engineering. AutoCkt [13] is a sparse sub-sampling RL technique optimizing the circuit parameters by taking discrete actions in the solution space. AutoCkt shows more efficiency over random RL agents and Differential Evolution. Still, it requires to be trained with thousands of SPICE simulations before deployment, which is costly.

In this paper we introduce DNN-Opt, a two-stage deep learning black-box optimization scheme, where we merge the strengths of Reinforcement Learning (RL), Bayesian Optimization (BO), and population-based techniques in a novel

978-1-6654-3274-0/21/$31.00

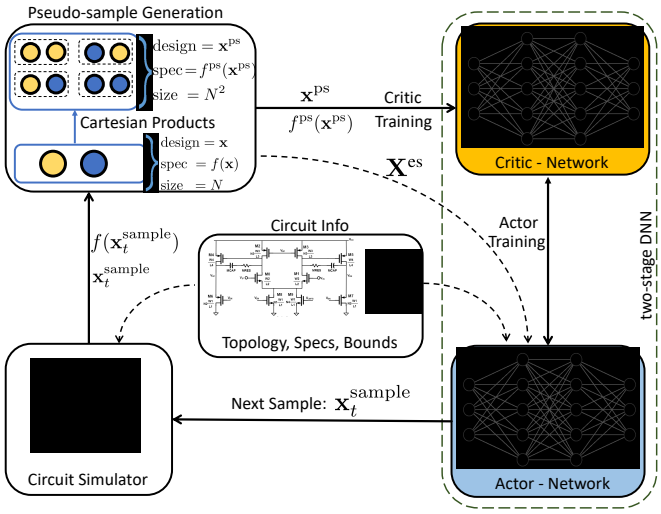

Fig. 1. DNN-Opt Framework

way. The key features of the DNN-Opt framework are below.

- We tailored a two-stage Deep Neural Network (DNN) architecture for black-box optimization tasks inspired by the actor-critic algorithms developed in the RL community.
- To leverage convergence behavior of population-based methods, DNN-Opt adopts a population-based search space control mechanism.
- We introduce a recipe for extending our work for large industrial designs using sensitivity analysis. In collaboration with a design house, we demonstrate that our work can also efficiently size large circuits with tens of thousands of devices in addition to small building blocks.

The rest of the paper is organized as follows. We formulate analog circuit sizing problem in Section II and introduce DNN-Opt with its RL core and other details. In Section III, the performance of DNN-Opt is demonstrated on small building blocks and large industrial circuits. We also provide performance comparisons of DNN-Opt with other optimization methods. The conclusions are provided in Section IV.

II. DNN-Opt Framework

A. Analog Circuit Sizing: Problem Formulation

We formulate analog circuit sizing task as a constrained optimization problem succinctly as below.

$$\begin{aligned} &\text{minimize} \;\; f_0(\mathbf{x}) \\ &\text{subject to} \;\; f_i(\mathbf{x}) \leq 0 \quad \text{for } i = 1, \ldots, m \end{aligned} \tag{1}$$

where, $\mathbf{x} \in \mathbb{D}^d$ is the parameter vector and d is the number of design variables of sizing task. Thus, $\mathbb{D}^d$ is the design space. $f_0(\mathbf{x})$ is the objective performance metric we aim to minimize. Without loss of generality, we denote i^{th} constraint by $f_i(\mathbf{x})$.

B. DNN-Opt Core: RL Inspired Two-Stage DNN Architecture

The overall framework of DNN-Opt is shown in Figure 1. DNN-Opt comprises a two-stage deep neural network architecture that interacts with a circuit simulator during the optimization process. The flow starts from generated samples in the design space; then, a critic-network is used to predict any new design point's performance. This prediction is used by the actor network to propose new candidates for simulation. This search scheme efficiently mimics BO behavior in space exploration. Besides, the sample generation is further optimized by adopting a population control scheme.

The two-stage network architecture of our work borrows its structure from Deep Deterministic Policy Gradient (DDPG) algorithm [14], which is an RL actor-critic algorithm [15] developed for continuous action spaces. However, actor-critic algorithms are not directly applicable to analog circuit sizing since it is not a Markov Decision Processes (MDP) [16], which is a *necessary condition* for any RL problem. Therefore we adapt DDPG algorithm with significant modifications tailored for analog circuit sizing.

In the context of analog circuit sizing, we will keep some of the RL notation but replace many for simplicity and clarity.
Design: A design is a set of circuit parameters which we denote by $\mathbf{x}$ and it is a vector of size d where each element corresponds to a particular design variable. The optimization goal is to find optimal $\mathbf{x}_{\text{opt}}$ which satisfies Eq. 1.
Population: A population is set of multiple designs.
Design Population Matrix: We define a design population matrix as $\mathbf{X} \in \mathbb{R}^{N \times d}$, where N is the population size. The parameters of i^{th} design is a row in the design population matrix $\mathbf{X}$, which is denoted as $\mathbf{x}_i$.
State Space: Our work maps optimization parameters (circuit design variables) to state representation in RL notation. A state of k^{th} design is transformed as $\mathbf{s}_k = \mathbf{x}_k$.
Action Space: Each action $\mathbf{a}_k$ in our new architecture corresponds to *change* in optimization parameters vector, $\mathbf{x}_k$, which can be denoted as $\mathbf{a}_k = \Delta\mathbf{x}_k$. An intuitive explanation of this choice is that an ideal action for an optimization task should propose change in each design variable to have a better design.
Critic-Network: Originally, a critic-network parameterized by θ^Q approximates the return value of an MDP $\text{Return} = Q(s_t, a_t|\theta^Q)$. We modify its role and use this network as a proxy in lieu of expensive SPICE simulator. Our modified critic-network provides a vector-to-vector mapping by taking an $(\mathbf{x}, \Delta\mathbf{x}) \in \mathbb{D}^{2d}$ as input and providing performance predictions $Q(\mathbf{x}, \Delta\mathbf{x}|\theta^Q) \in \mathbb{R}^{m+1}$ at output, one-dimension is for objective specification and m for constraint specifications.
Actor-Network: An actor-network parameterized by θ^μ would take a state as its input and determine an action to take $\mathbf{a}_k = \mu(\mathbf{s}_k|\theta^\mu)$. In the context of analog circuit sizing, actor-network provides change in design parameter vector for design k as: $\Delta\mathbf{x}_k = \mathbf{a}_k = \mu(\mathbf{x}_k|\theta^\mu)$.
Critic-Network Training: We utilize critic-network for modeling design variable to circuit performance relationship. For effective training, we use data augmentation techniques to generate N^2 *pseudo-samples* (ps) using original N samples. In order to generate pseudo-samples, we use two-samples $\mathbf{x}_i$ and $\mathbf{x}_j$ and corresponding spec vectors $f(\mathbf{x}_i)$ and $f(\mathbf{x}_j)$, as follows:

$$\begin{aligned} \mathbf{x}_{ij}^{\text{ps}} &= [\mathbf{x}_i, \Delta\mathbf{x}_{ij}] = [\mathbf{x}_i, \mathbf{x}_j - \mathbf{x}_i] \\ f^{\text{ps}}(\mathbf{x}_{ij}^{\text{ps}}) &= f(\mathbf{x}_j) \end{aligned} \tag{2}$$

This leads to change in the input dimensionality of critic-network from d to $2d$ since we now have to use $(\mathbf{x}, \Delta\mathbf{x})$ instead of $\mathbf{x}$ or $(\mathbf{x}+\Delta\mathbf{x})$. Our experiments conducted on Bayesmark [17] benchmark problems showed that using $2d$ inputs and training with pseudo-samples boosted critic-network's accuracy significantly over a network trained with d inputs and original samples.

For a batch-size of N_b pseudo-samples, the following Mean Squared Error (MSE) loss function is used to train the critic network.

$$\mathrm{L}(\theta^Q) = \frac{1}{N_b(m+1)} \sum_{k=1}^{N_b} \sum_{l=1}^{m+1} \left(Q(\mathbf{x}_k, \Delta\mathbf{x}_k)^l - f(\mathbf{x}_k + \Delta\mathbf{x}_k)^l\right)^2 \quad (3)$$

where $Q(\mathbf{x}_k, \Delta\mathbf{x}_k)^l$ is the critic-network's approximation for k^{th} pseudo-sample's l^{th} performance and $f(\mathbf{x}_k + \Delta\mathbf{x}_k)^l$ is the SPICE simulated value for the same design-performance pair. To clarify, we have SPICE simulation values for pseudo-samples because the way they are constructed.

Actor-Network Training: Training of actor-network is done after critic-network is trained and its hyperparameters are fixed. The training of actor-network corresponds to search in design space for *better* designs. We come up with a Figure of Merit (FoM) function, $g(\cdot)$, based on performance-vector to objectively quantify how better a design is with respect to others.

$$g\left[f(\mathbf{x})\right] = w_0 \times f_0(\mathbf{x}) + \sum_{i=1}^{m} \min\left(1, \max(0, w_i \times f_i(\mathbf{x}))\right) \quad (4)$$

where w_i is the weighting factor. Note, a $\max(\cdot)$ clipping used for equating designs after constraint are met and $\min(\cdot)$ clipping is used for practical purposes to prevent single constraint violation to dominate $g(\cdot)$ value. We train actor-network parameters by using $g(\cdot)$ function and replacing SPICE simulation values $f(\cdot)$ by the critic-network predictions $Q(\mathbf{x}, \Delta\mathbf{x})$. We will further use a population of "elite" solutions (es) of size N_{es} to restrict search space for actor network. Population of elite solutions is a subset of total population determined based on the FoM ranking.

For a batch-size of N_b samples the following loss-function is used to train actor network.

$$L\left(\theta^\mu\right) = \frac{1}{N_b} \sum_{k=1}^{N_b} \left(g\left[Q(\mathbf{x}_k, \mu(\mathbf{x}_k \mid \theta^\mu))\right] + \|\lambda * \mathrm{viol}_k\|_2\right) \quad (5)$$

where $\mu(\mathbf{x}_k \mid \theta^\mu)$ is proposed parameter change vector $\Delta\mathbf{x}_k$ by the actor network. $(\lambda * \mathrm{viol}_k)$ is an element-wise vector multiplication where λ is weighting coefficient chosen to be very large to prevent any boundary violation and keep the search in the restricted search region. The total boundary violation viol_k for action k is defined as follows:

$$\mathrm{viol}_k = \max(0, lb_{\text{rest}} - (\mathbf{x}_k + \Delta\mathbf{x}_k)) + \max(0, (\mathbf{x}_k + \Delta\mathbf{x}_k) - ub_{\text{rest}}) \quad (6)$$

where lb_{rest} and ub_{rest} are the restriction boundary vectors for design variables determined by the population of elite solutions given by:

$$lb^i_{\text{rest}} = \min(\mathbf{x}^i) \;\; \forall i = 1, \ldots, d$$
$$ub^i_{\text{rest}} = \max(\mathbf{x}^i) \;\; \forall i = 1, \ldots, d$$

where, $\mathbf{x}^i$ is the column vector of size N_{es} consisting of i^{th} parameter of all designs in the elite population.

The hyperparameters (number of layers, number of nodes, learning rate, etc.) of the architecture for the actor and critic networks were found based on empirical studies.

C. Sensitivity Analysis

We use sensitivity analysis to prune design search space for efficiently finding an optimized solution. A blind search space exploration may lead to wasted circuit simulations during optimization. For example, in a classical seven transistor Operational Amplifier (OpAmp) [4] power dissipation does not depend on the differential pair devices once they are in saturation. Thus, if we want to size a circuit for reducing power, we should not make device properties of the differential pair devices as variables. To use sensitivity analysis in practice for any generic circuit, we first traverse the circuit hierarchy and collect all unique device design variables, d. Then, we perform sensitivity analysis by perturbing each of the design variables around its nominal value and observing its impact on objective and constraints, f_i. More formally, we compute sensitivity $\mathcal{S}_{ij}$ as

$$\mathcal{S}_{ij} = \frac{\delta f_i}{\delta d_j}, \forall i = 0, \ldots, m; j = 1, \ldots, d. \quad (7)$$

We only need to consider design variables for which $\mathcal{S}_{ij} > thresh$, where $thresh$ is a user-defined number. Empirically, this analysis prunes design search space effectively, allowing us to work on large scale circuits.

We are now ready to present the overall framework of DNN-Opt in the next subsection.

D. DNN-Opt: Overall Framework

The overall framework for DNN-Opt is provided in Algorithm 1. As a prerequisite, we apply sensitivity analysis for a large design and reduce number of design variables to a workable range. We then randomly sample N_{init} points from the design search space to build initial population. For optimization iteration t, first step is to initialize actor-critic parameters followed by pseudo-sample generation. Next actor-network and critic-network are trained. After this, an elite-population is constructed based on FoM of total-population (this elite-population will be updated with optimization iterations). The next query point is generated from elite-population, $\mathbf{X}^{\text{es}}$, using pre-trained actor-critic as follows. We use every design, $\mathbf{x}^{\text{es}}_i$, in the pool of elite-population as input to actor-network. The output of actor-network, $\Delta\mathbf{x}^{\text{es}}_i = \mu(\mathbf{x}^{\text{es}}_i)$, is proposed change for design parameters in search of an optimal solution. With the imposed exploration noise $(\mathcal{N})$, a candidate design point is naturally formed as: $\mathbf{x}^{\text{ca}}_i = \mathbf{x}^{\text{es}}_i + \mu(\mathbf{x}^{es}_i) + \mathcal{N}$. At this step, we have exactly the same number of proposed candidates, $\mathbf{X}^{\text{ca}} = [\mathbf{x}^{\text{ca}}_i, \ldots, \mathbf{x}^{\text{ca}}_{N_{\text{es}}}]$, as the size of elite-population. Once the population pairs, $\mathbf{X}^{\text{es}}$ and $\mathbf{X}^{\text{ca}}$, are formed the next sample point for iteration t is selected using Eq. 8.

$$\mathrm{x}^{\text{sample}}_t = \left[\mathbf{x}^{\text{ca}}_k \text{ for } k = arg\,min_i \left(g[Q(\mathbf{x}^{\text{es}}_i, \mathbf{x}^{\text{ca}}_i - \mathbf{x}^{\text{es}}_i)]\right)\right] \quad (8)$$

Algorithm 1 DNN-Opt Algorithm

Require: Dimensionality reduction with sensitivity analysis **if** design is *large*

Require: An initial sample set $\mathbf{X}^{\text{init}}$ of N_{init} designs and their evaluations $f(\mathbf{X}^{\text{init}})$

1: Define total population $\mathbf{X}^{\text{tot}} = \mathbf{X}^{\text{init}}$
2: **for** $t = 1, 2, \ldots, t_{max}$ **do**
3: Initialize actor & critic network parameters θ^{μ} and θ^{Q}
4: Generate pseudo-samples using existing design $\mathbf{X}^{\text{tot}} \rightarrow$ Eqn. 2
5: Train critic-network $\rightarrow$ Eqn. 3
6: Train actor-network $\rightarrow$ Eqn. 5
7: Calculate FoM for each design by FoM $= g[f(\mathbf{X}^{\text{tot}})]$
8: Choose N_{es} designs with smallest FoM to form population of elite solutions $\mathbf{X}^{\text{es}}$.
9: Find query point (next sample) $\mathbf{x}_t^{\text{sample}}$ using actor-model $\rightarrow$ Eqn. 8
10: Simulate the query point and obtain specs $f(\mathbf{x}_t^{\text{sample}})$ via SPICE sims
11: **if** return cond(e.g. specs are met) **then**
12: break
13: **end if**
14: $\mathbf{X}^{\text{tot}}$.append($\mathbf{x}_t^{\text{sample}}$)
15: Go back to line 3
16: **end for**
17: **return** The design with highest FoM

III. Experimental Results

To demonstrate the reliability and efficiency of the DNN-Opt, we apply it to two sets of experiments using six circuit examples. The first experiment set is on small building blocks where every transistor is parameterized and sized, and the second experiment set includes larger industrial circuits with thousands of nodes and devices.

A. Experiments with Small Building Blocks

We tested DNN-Opt on two small building blocks: a folded cascode amplifier and a strong-arm latch comparator. We included the majority of the circuit performances in the constraint list to mimic real-world design experience. Both designs are implemented in 180nm CMOS technology.

We compare our algorithm with three other well-known methods: a) A Differential Evolution (DE) method, which is a conventional population-based model-free algorithm, b) Bayesian Optimization with weighted Expected Improvement (BO-wEI) [11], which is a modified version of Bayesian Optimization for constrained problems, and c) GASPAD method [9], a surrogate model (GP) assisted evolutionary framework. To account for the randomized techniques involved in all these methods, we repeat experiments ten times to report each method's findings. We determine the simulation budgets for our experiments by considering the convergence nature of the methods. DE has a simulation budget of 10000, and BO-wEI, GASPAD, and DNN-Opt are limited by 500 simulations. All the experiments are run on a workstation with Intel Xeon CPU and 128GB RAM, and a commercial SPICE simulator. We

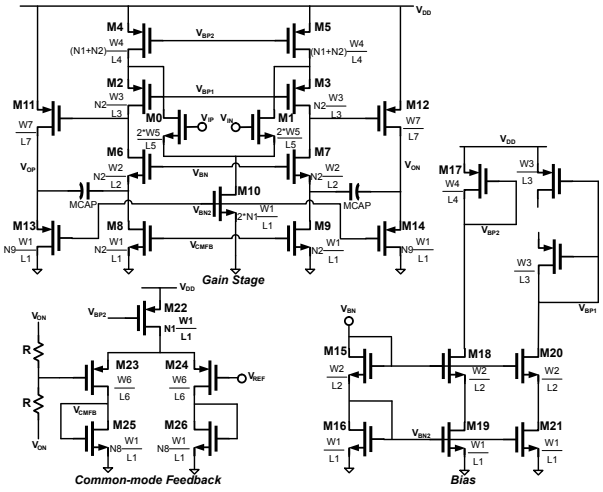

Fig. 2. Schematic of the folded-cascode OTA

used several metrics to compare the algorithms. We provide statistics of the methods for each example, and we denote the number of times a feasible solution is found by *success rate*. We also share the evolution of FoM value calculated based on Eq. 4 to demonstrate each algorithm's convergence during runtime. The constraint expressions given in Eq. 9 and 10 can be trivially readjusted to fit into the form of Eq. 1.

Folded Cascode OTA: The first test case is a two-stage folded-cascode Operational Transconductance Amplifier (OTA) (Figure 2).It has 20 design variables, and the designer provided search ranges are as shown in Table I.

TABLE I
DESIGN PARAMETERS AND RANGES FOR THE FOLDED-CASCODE OTA

Parameter Name	Unit	LB	UB
L1-L2-L3-L4-L5-L6-L7	μm	0.18	2
W1-W2-W3-W4-W5-W6-W7	μm	0.24	150
N1-N2-N8-N9	integer	1	20
MCAP	fF	100	2000
Cf	fF	100	10000

W:device width; L:device length; UB:upper bound; LB:lower bound

The sizing problem is defined as follows:

$$\begin{aligned} &\text{minimize Power} \\ &\text{s.t.} \quad \text{DC Gain} > 60 \text{ dB} \qquad \text{Settling Time} < 30 \text{ ns} \\ &\qquad \text{CMRR} > 80 \text{ dB} \qquad \text{Saturation Margin} > 50 \text{ mV} \\ &\qquad \text{PSRR} > 80 \text{ dB} \qquad \text{Unity Gain Freq.} > 30 \text{ MHz} \\ &\qquad \text{Out. Swing} > 2.4 \text{ V} \qquad \text{Out. Noise} < 30 \text{ mV}_{\text{rms}} \\ &\qquad \text{Static error} < 0.1 \qquad \text{Phase Margin} > 60 \text{ deg.} \end{aligned} \tag{9}$$

In our experiment, the following transistors are required to operate in the saturation region: M1, M3, M4, M7, M9, M10, M12, M13, and [M15-M26]. The total number of design constraints becomes 29.

The statistical results for all the reference algorithms are shown in Table II. DNN-Opt shows high reliability and find a feasible solution in all its trials. However, other model-based methods, BO-wEI and GASPAD, fail to achieve similar behavior. DE can also find feasible results, but DNN-Opt is 24x more efficient in the number of required simulations to find the first feasible result. It is also demonstrated in Table II that, on average, the final design proposed by DNN-Opt

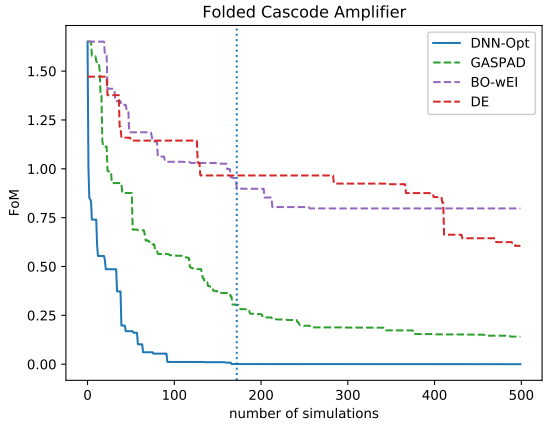

Fig. 3. The average FoM (lower is better) curve for 500 simulations

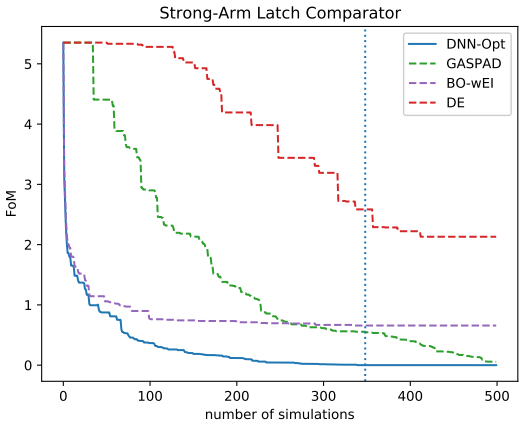

Fig. 4. The average FoM (lower is better) curve for 500 simulations

TABLE II
STATISTICS FOR DIFFERENT ALGORITHMS: FOLDED CASCODE OTA

Algorithm	DE	BO-wEI	GASPAD	**DNN-Opt**
success rate	10/10	2/10	4/10	**10/10**
# of simulations	3200	>500	>500	**132**
Min power (mW)	0.75	0.91	0.72	**0.62**
Max power (mW)	1.53	1.62	1.75	**0.77**
Mean power (mW)	1.14	1.25	0.96	**0.71**
Modeling time (h)	NA	30	6.5	**0.6**
Simulation time (h)	54	2.7	2.7	2.7
Total runtime (h)	54	32.7	8.2	**3.3**

draws up to 43% less power. The modeling time required by DNN-Opt is up to 50x smaller compared to other model-based methods. This results in 2.5–16x efficiency for total runtime.

Figure 3 includes the FoM curve with iterations, where DNN-Opt shows strong convergence behavior and outperforms other methods. For our ten runs, DNN-Opt finds the feasible solution within 205 iterations (marked with vertical dashed line) across all its ten trials. Although it is slow, GASPAD shows convergence to optimal FoM, but we observed that BO-wEI is often trapped in local optima.

Strong-Arm Latch Comparator: The second test case is SA-Latch Comparator, which is shown in Figure 5. It has 13 design variables, and their names and bounds are shown in table III.

TABLE III
DESIGN PARAMETERS AND THEIR RANGES FOR SA-LATCH COMPARATOR

Parameter Name	Unit	LB	UB
L1-L2-L3-L4-L5-L6	μm	0.18	10
W1-W2-W3-W4-W5-W6	μm	0.22	50
CL_finger	integer	10	300

The constrained optimization problem consists of 10 constraints in total:

$$
\begin{aligned}
&\text{minimize Power} \\
&\text{s.t.} \quad \text{Set Delay} < 10 \text{ ns} \\
&\qquad \text{Reset Delay} < 6.5 \text{ ns} \\
&\qquad \text{Area} < 26 \ \mu\text{m}^2 \\
&\qquad \text{Input-referred Noise} < 50 \ \mu\text{Vrms} \\
&\qquad \text{Differential Reset Voltage} < 1 \ \mu\text{V} \\
&\qquad \text{Differential Set Voltage} > 1.195 \text{ V} \\
&\qquad \text{Positive-Integration Node Reset Voltage} < 60 \ \mu\text{V} \\
&\qquad \text{Negative-Integration Node Reset Voltage} < 60 \ \mu\text{V} \\
&\qquad \text{Positive-Output Node Reset Voltage} < 0.35 \ \mu\text{V} \\
&\qquad \text{Negative-Output Node Reset Voltage} < 0.35 \ \mu\text{V}.
\end{aligned} \tag{10}
$$

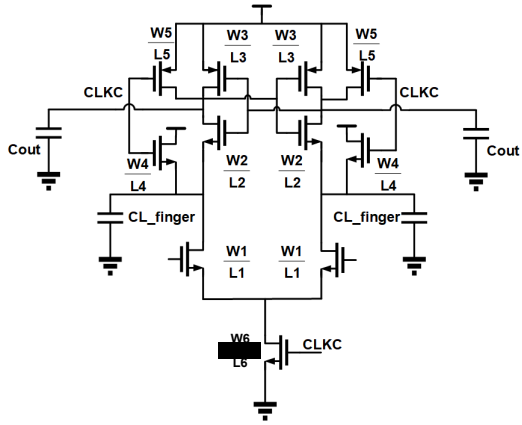

Fig. 5. Schematic of SA-Latch Comparator

The statistical results for all the reference algorithms are shown in Table-IV. Due to relatively tighter constraints for SA-Latch Comparator, methods typically needed a larger number of simulations to converge. DDN-Opt is the only method that finds a feasible solution in all trials, and our method shows more than 30x efficiency compared to DE. GASPAD shows relatively competitive results, but DNN-Opt finds a solution with 25% better power consumption than successful runs of GASPAD. The runtime observations are similar to the folded cascode case.

FoM curves are shown in Figure 4 for different methods. DNN-Opt finds a feasible solution within 348 simulations, which is much earlier than the others. BO-wEI shows a similar convergence trend for initial iterations then fails to model one of the constraints properly. Our observations showed that all the runs with the BO-wEI method were unable to meet input-referred noise, and some failed for set delay.

TABLE IV
SA LATCH COMPARATOR RESULTS

Algorithm	DE	BO-wEI	GASPAD	DNN-Opt
success rate	5/10	0/10	6/10	**10/10**
# of simulations	>10000	>500	>500	**330**
min (μW)	2.98	NA	3.05	**2.50**
max (μW)	4.22	NA	3.75	**2.75**
mean (μW)	3.57	NA	3.45	**2.65**
Modeling time (h)	NA	17	3	**0.3**
Simulation time (h)	72	3.6	3.6	3.6
Total runtime (h)	72	20.6	6.6	**3.9**

B. Experiments with Industrial Scale Circuits

We tested DNN-Opt on four industrial circuits designed at a very advanced technology node. These circuits were already in the process of manual sizing by expert analog designers and needed some fine-tuning. For these industrial circuits, we did not have access to other algorithms (DE, GASPAD, BO-wEI), and hence our baseline is with a commercial black-box optimizer based on Simulated Annealing. As will be demonstrated in this section, DNN-Opt performs well on large circuits and is not limited to small examples. Analog designers assisted in selecting permissible parameter ranges of the devices, considering layout impacts and process rules. For industrial cases, we identify critical devices based on Eq. 7 for the failing constraints (f_i's of Eq. 1). Note, MLParest [18] was used in the loop of DNN-Opt which helps analog designer estimate post-layout effects early in the design.

Inverter Chain: The first case is a simple inverter chain used mainly for tool development and flow testing. We used all the devices (8) in the four stage inverter chain. There were only two specs, delay and power.

Level Shifter: Sensitivity analysis identified ten critical devices impacting failing performances, and that led to a design space of 3.9×10^{15}. There were 60 total specs like delay, rise, fall, power, current, etc.

Low-Dropout (LDO) Regulator: We used sensitivity analysis to identify six critical devices leading to search space of 1.6×10^{13}. The circuit had PSRR, Gain Margin, Phase Margin, DC Gain, GBW, etc., as part of nine constraints. The number of devices is high due to arrayed instances used by the analog engineer.

Continuous-Time Linear Equalizer (CTLE): Sensitivity analysis identified eight critical devices impacting failing performances. With design parameter and ranges identified by analog designers, we had a design space of 3.3×10^{25}. There were a total of 14 constraints like DC Gain, offset, Nyquist Gain, Fpeak, Peaking Max, Power, etc.

As illustrated in Table-V, DNN-Opt outperforms commercial optimizer available in the industry in terms of the number of simulations required to meet the constraints by 5x. We would like to emphasize that we can deal with fairly complex CTLE circuit by using 4x smaller number of costly SPICE simulations. Additionally, the optimal solution proposed by DNN-Opt consumed 8% lesser power than simulated annealing. Our examples represent real use cases where designers already spend several days worth of human time in fixing constraints. Had we started with designs without any knowledge of human designers baked-in, we would have seen even greater returns in sample efficiency like Section III-A.

TABLE V
DNN-OPT RESULTS ON INDUSTRIAL CIRCUITS

Circuit	MOS	Nodes	Simulated Annealing (SA)	**DNN-Opt**
Inverter Chain	8	7	>1000	**90**
Level Shifter	1.2k	3.9k	1200	**195**
LDO	167k	2.8k	552	**112**
CTLE	173k	63k	587	**150**

Number of SPICE simulations shown in column SA and DNN-Opt for meeting constraints (lower is better).

IV. CONCLUSION

In this work, we presented DNN-Opt, a novel sample efficient black-box optimization algorithm that combined the strengths of deep neural networks and reinforcement learning paradigm. We also give a recipe to extend our work for large circuits with thousands of devices. Our algorithm's effectiveness has been successfully demonstrated on various circuit building blocks and large industrial circuits leading to 5–30x sample efficiency, while being able to find feasible solution for all circuit sizing tasks and showing superior converge curves compared to other methods.

ACKNOWLEDGEMENT

This work is supported in part by NSF under Grant No. 1704758.

REFERENCES

[1] N. Horta, "Analogue and mixed-signal systems topologies exploration using symbolic methods," *Analog Integr. Circuits Signal Process.*, 2002.

[2] N. Jangkrajarng, S. Bhattacharya, R. Hartono, and C.-J. Shi, "Iprail—intellectual property reuse-based analog ic layout automation," *Integration*, 2003, analog and Mixed-signal IC Design and Design Methodologies.

[3] W. Daems, G. Gielen, and W. Sansen, "Simulation-based generation of posynomial performance models for the sizing of analog integrated circuits," *IEEE TCAD*, 2003.

[4] M. d. Hershenson, S. P. Boyd, and T. H. Lee, "Optimal design of a cmos op-amp via geometric programming," *IEEE TCAD*, 2001.

[5] Y. Wang, M. Orshansky, and C. Caramanis, "Enabling efficient analog synthesis by coupling sparse regression and polynomial optimization," in *DAC*, 2014.

[6] R. Acar Vural and T. Yildirim, "Analog circuit sizing via swarm intelligence," *AEU - International Journal of Electronics and Communications*, 2012.

[7] B. Liu, G. Gielen, and F. V. Fernndez, *Automated Design of Analog and High-frequency Circuits: A Computational Intelligence Approach*. Springer, 2013.

[8] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[9] B. Liu, D. Zhao, P. Reynaert, and G. G. E. Gielen, "Gaspad: A general and efficient mm-wave integrated circuit synthesis method based on surrogate model assisted evolutionary algorithm," *IEEE TCAD*, Feb 2014.

[10] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *NIPS, 2012*.

[11] W. Lyu, F. Yang, C. Yan, D. Zhou, and X. Zeng, "Multi-objective bayesian optimization for analog/rf circuit synthesis," in *DAC, 18*.

[12] H. Wang, K. Wang, J. Yang, L. Shen, N. Sun, H. Lee, and S. Han, "Gcn-rl circuit designer: Transferable transistor sizing with graph neural networks and reinforcement learning," *DAC*, 2020.

[13] K. Settaluri, A. Haj-Ali, Q. Huang, K. Hakhamaneshi, and B. Nikolić, "Autockt: Deep reinforcement learning of analog circuit designs," *DATE*, 2020.

[14] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning." in *ICLR*, 2016.

[15] V. Konda and J. Tsitsiklis, "Actor-critic algorithms," in *SIAM Journal on Control and Optimization*. MIT Press, 2000.

[16] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: A Bradford Book, 2018.

[17] R. Turner and D. Eriksson, "Bayesmark," https://github.com/uber/bayesmark, 2020.

[18] B. Shook, P. Bhansali, C. Kashyap, C. Amin, and S. Joshi, "Mlparest: Machine learning based parasitic estimation for custom circuit design," in *DAC*, 2020.