

Provenance-Aware Physics-Guided Feasibility-Region Learning for Pre-layout Two-Stage CMOS Op-Amp Sizing: A Calibration-Ready Workflow

Zihan Wang

*Glasgow College, University of Electronic Science and Technology of China, Chengdu, China
2023190904031@std.uestc.edu.cn*

Abstract. Analog operational-amplifier sizing is a costly multi-objective search problem: transistor dimensions and bias conditions jointly determine gain, bandwidth, stability, slew rate, rejection ratios, common-mode range, power, and offset. This paper presents a provenance-aware, physics-guided feasibility-region learning workflow for a two-stage CMOS operational amplifier. The report-anchored working set contains one TT baseline and 1,399 grouped-7D schematic-level rows. An independent LTspice 17.2.4 experiment contributes 96 portable schematic-level rows, including 32 complete cases, of which 20 satisfy the portable feasibility criteria. Six algorithmic entries are evaluated on the complete cases: Direct Ridge, physics-augmented Ridge, direct ExtraTrees, physics-augmented ExtraTrees, multilayer perceptron (MLP), and a constrained GP-UCB-style acquisition reference. The portable audit reports a mean normalized mean absolute error (MAE) of 0.393 for the policy reference, while MLP obtains the highest feasibility area under the precision-recall curve (AUC-PR) of 0.980. A one-round offline active-sampling replay ranks 5,000 grouped-7D candidates and selects 20 points; the macro normalized MAE remains 0.0664 before and after both the active and equal-size random additions. The report anchor is TT/pre-layout, and the working-set and portable-audit results are interpreted within their declared provenance rather than as process-general performance claims.

Keywords: two-stage CMOS operational amplifier, physics-guided feasibility-region learning, analog operational-amplifier sizing

1. Introduction

Analog CMOS operational amplifiers are fundamental building blocks in mixed-signal and analog integrated circuits, supporting gain, buffering, filtering, and feedback. Their electrical performance affects system-level accuracy, stability, energy efficiency, and development time. Nevertheless, sizing a two-stage CMOS operational amplifier remains difficult because device dimensions and bias conditions are strongly coupled, while the usable design region is constrained by multiple small-signal, large-signal, and operating-region requirements. An efficient design-automation method is therefore needed to identify promising regions before expensive circuit-level verification.

Recent learning-based methods have explored analog sizing through deep reinforcement learning and graph representations. AutoCkt and GCN-RL formulate transistor sizing as learned search using reinforcement learning and graph representations [1, 2]. CALT treats sizing as classification with adaptive labeling thresholds [3]. Domain-knowledge-infused learning injects circuit knowledge into analog/RF optimization, while MA-Opt uses multiple actors to explore alternative designs in parallel [4, 5].

tSS-BO, circuit-attention actor–critic learning, and procedural or sparse-reward reinforcement learning address high-dimensional or robust analog sizing through subspace sampling, structured representations, electrical design spaces, or procedural abstractions [6-8]. High-dimensional Bayesian optimization, device-representation learning, multi-objective search, and derivative-aware acquisition improve sample efficiency and trade-off exploration [9-12]. Together, these works improve search efficiency, but they generally do not make the provenance of mixed report-derived, working, and simulator-backed labels explicit, treat the multi-metric feasible boundary as a first-class learning target, or specify a reproducible replacement path from a report-derived baseline to process-specific PSF data.

The research gap leads to two questions: RQ1, whether operating-region priors improve prediction or calibration over direct design variables; and RQ2, whether boundary-aware acquisition improves feasible-region discovery under a matched label budget. The process-specific PSF replacement is defined as a calibration protocol rather than evaluated as a third research question in this paper. The contributions are fivefold: a provenance-aware data schema for report, working, and simulator-derived rows; a physics-guided surrogate using low-cost operating-region features; a boundary- and novelty-aware active-sampling rule; an explicit TT/pre-layout calibration protocol; and a portable simulator-backed comparison against algorithmic baselines. The remainder is organized as follows. Section 2 defines the circuit, constraints, physics features, surrogate, and acquisition rule. Section 3 describes dataset construction, measurement definitions, and evaluation protocol. Section 4 presents the results, and Section 5 concludes the paper.

2. Problem formulation and the proposed physics-guided method

2.1. Circuit and design variables

The target is a two-stage Miller-compensated CMOS operational amplifier in the ST02 0.5- μm process. The report anchor is TT, 27 °C, pre-layout schematic level, with $V_{DD}=5\text{ V}$ and an external output load $C_L = 3\text{ pF}$. The physical schematic contains the internal compensation network $R_2 = 5.385\text{ k}\Omega$ and $C_2 = 3.042\text{ pF}$; C_2 and C_L are separate elements. The design vector is

$$x = [W_{MP12}, W_{MP3}, W_{MN12}, W_{MN34}, W_{MN5}, R_2, C_2]T \quad (1)$$

Where W_{MP12} denotes the common width of MP1 and MP2; W_{MP3} denotes MP3; W_{MN12} denotes the common width of MN1 and MN2; W_{MN34} denotes the common width of MN3 and MN4; W_{MN5} denotes MN5; R_2 is the series nulling-resistor value; and C_2 is the internal Miller-compensation capacitance.

The channel length is fixed at $L=2.61\text{ }\mu\text{m}$ for MP1–MP3 and MN1–MN5, and the bias current is fixed at $I_{bias}=20\text{ }\mu\text{A}$ for the nominal report anchor. Prior automated analog-sizing studies likewise use transistor dimensions and bias-related quantities as design coordinates [1, 2].

The parameter bounds are explicit in the grouped-7D dataset: each width, R_2 , and C_2 is varied by $\pm 50\%$ around the confirmed schematic anchor, while $L = 2.61 \mu\text{m}$ and $I_{bias} = 20 \mu\text{A}$ remain fixed. The grouped variables map to MP1/MP2: $W = 11 \mu\text{m}$; MP3: $W = 208.8 \mu\text{m}$; MN1/MN2: $W = 18.27 \mu\text{m}$; MN3/MN4: $W = 2.61 \mu\text{m}$; and MN5: $W = 28.51 \mu\text{m}$. The report source excitation and all metric definitions are fixed in Section 3.2; no individual transistor-length or bias-current variable is included as an independent coordinate.

The grouped parameterization is used consistently by the model, tables, figures, and nominal measurement plan. The seven grouped variables are the only learned design coordinates.

2.2. Multi-metric feasibility

The performance vector is given below.

$$y = [A_{vd}, GBW, PM, SR, CMRR, PSRR, ICMR_L, ICMR_H, P_{total}, VOS_{proxy}]^T \quad (2)$$

Where A_{vd} is the differential voltage gain; GBW, PM, and SR denote gain-bandwidth product, phase margin, and slew rate, respectively. CMRR and PSRR are the common-mode and power-supply rejection ratios; $ICMR_L$ and $ICMR_H$ are the lower and upper input common-mode limits; P_{total} is total power; and VOS_{proxy} is the input-offset proxy.

Table 1. Multi-metric specifications and report baseline

Metric	Req.	Unit	Report
Gain	≥ 70	dB	70.197
GBW	≥ 10	MHz	20.190
PM	≥ 60	deg	74.895
SR	≥ 10	V/ μs	20.190
CMRR	≥ 80	dB	92.641
PSRR	≥ 80	dB	89.740
ICMR-L	≤ 2	V	1.500
ICMR-H	≥ 4	V	4.400
Power	≤ 2	mW	1.261
VOS linearity	diagnostic	mV; V	linear trace; 1.200 mV

The thresholds and report-anchor values in Table 1 define the scalar margins used below. Let g_k (y) denote the signed margin of scalar constraint k. The current working feasibility flag is defined by

$$\begin{cases} F(x) = 1 & \Leftrightarrow \forall k \in \{1, \dots, 9\} \\ & g_k(y_k(x)) \geq 0 \end{cases} \quad (3)$$

Where $F(x)$ is the feasibility flag, g_k is the signed margin of constraint k , $y(x)$ is the metric vector, and k indexes the nine scalar specifications.

In Eq. (3), VOS linearity is retained as a separate diagnostic because the current working rows store VOS_{proxy} rather than a row-level fitted residual. The all-constraints feasibility convention follows specification-driven analog sizing and feasibility-region exploration [3, 6].

For implementation, all margins use a common direction. For a lower-bound metric $y_k \geq \tau_k$, $g_k = y_k - \tau_k$; for an upper-bound metric $y_k \leq \tau_k$, $g_k = \tau_k - y_k$. The normalized minimum margin is

$$\begin{cases} m_{min}(x) = \frac{\min_k g_k(y(x))}{s_k} \\ s_k = \max(|\tau_k|, \varepsilon) \end{cases} \quad (4)$$

2.3. Operating-region features

The direct surrogate uses only x . The physics-guided surrogate augments x with nine features: a gm/ID proxy, a $gmro$ proxy, four transistor saturation margins, a headroom proxy, a saturation-region ratio, and a current-density proxy. For a future real run, gm/ID should be computed from OP values, $gmro$ from $gm-ro$ or an equivalent small-signal quantity, and saturation margins from terminal voltages and $VDSAT$. The current working data use these features to test the interface; they are not claimed to be measured OP values. The use of gm/ID as an operating-region descriptor is motivated by established design practice for CMOS two-stage operational amplifiers [13].

$$\rho_{sat} = 1/4 \sum_{i=1..4} \sigma(\alpha m_i) \quad (5)$$

Where σ is the logistic function.

2.4. Physics-guided surrogate model

For each metric, a quadratic Ridge model is trained. The direct model uses the seven grouped variables, whereas the physics-guided model appends nine operating-region features. Eight bootstrap fits provide a prediction mean and a simple uncertainty estimate. Low-order surrogate models also provide a computationally efficient baseline for simulation-based analog sizing [6, 9].

2.5. Uncertainty-aware active sampling

To prioritize simulations that are both informative and close to the specification boundary, the acquisition score combines four signals: ensemble uncertainty, boundary proximity, predicted feasibility, and design-space novelty. The score is evaluated for every candidate before the next simulation batch is selected.

$$S(x) = 0.55U(x) + 0.25B(x) + 0.10(1 - P_{feas}(x)) + 0.10N(x) \quad (6)$$

Where U is the clipped ensemble uncertainty, B is boundary proximity, predicted feasibility probability is P_{feas} , and N is standardized novelty.

The active score is defined as

$$\begin{cases} B = \exp(-|m_{min}|) \\ N = \min(1, d_{std}/4) \end{cases} \quad (7)$$

The score uses U normalized by the training-set output spread and P_{feas} as the mean bootstrap feasibility probability. Candidates are ranked by S and filtered by a minimum standardized distance before 20 points are selected. The combination of uncertainty, boundary proximity, and novelty is related to active-learning feasibility exploration and Bayesian acquisition for analog sizing [6, 12].

The workflow maps the design vector to circuit-level performance metrics under fixed nominal conditions, appends operating-region features, and ranks candidates using uncertainty, specification-boundary proximity, predicted feasibility, and novelty within the simulation budget.

3. Dataset construction and evaluation protocol

3.1. Data sources and provenance

There are 280 test points and 1,120 training points in the working dataset; the test and training point distribution is fixed based on a provided seed. All the rows are based on the confirmed physical schema grouped-7D, while $L=2.61 \mu\text{m}$ and $I_{bias}=20 \mu\text{A}$ are fixed parameters, and not independent variables. One baseline row uses the report values directly, and 1,399 rows are grouped-7D working rows generated around that anchor. VOS_{proxy} is calibrated from the report's near-linear VOS -versus- VCM figure and its 1.2-mV TT offset anchor. The curve images are not stored or used as scalar measurements, but as visual reference.

The report-derived row is the nominal TT/pre-layout experimental anchor. The grouped-7D expansion is used to evaluate the learning and sampling pipeline, and the independent LTspice 17.2.4 track provides 96 portable schematic-level rows: 32 complete rows with all ten output labels and 64 partial rows. Of the complete rows, 20 satisfy the portable feasibility criteria. This track is used only as an algorithm-validation reference and does not replace ST02-device evidence.

3.2. Measurement definitions and nominal TT/pre-layout conditions

The nominal measurement contract uses the ST02 0.5- μm PDK at the TT process corner, 27 °C, $VDD = 5 \text{ V}$, pre-layout schematic level, and an external $C_L = 3 \text{ pF}$. The schematic includes the internal compensation elements $R_2 = 5.385 \text{ k}\Omega$ and $C_2 = 3.042 \text{ pF}$; C_L is separate from C_2 . In the open-loop AC test, both inputs are biased at 2.5 V DC, while V_{IP} and V_{IN} receive +1 V and -1 V AC phasors, respectively. The differential frequency response is swept from 1 Hz to 1 GHz. These phasors define the linear AC test and are not the large-signal transient step used to measure SR. Each row stores the analysis name, f_{ref} for CMRR/PSRR, AC source settings, expression version, and convergence flag. The SR record stores the transient-step amplitude, edge time, simulation duration, output swing, and measurement window; the ICMR record stores the common-mode sweep step and output-error tolerance; and the VOS record stores the fit interval, R^2 , and maximum residual. The same settings are used for all compared rows.

The differential gain is defined as.

$$A_{vd}(f) = |V_{OUT}(f)/(V_{IP}(f) - V_{IN}(f))| \quad (8)$$

These gain, bandwidth, and stability definitions follow standard CMOS operational-amplifier characterization practice [13].

The unity-gain frequency and phase margin are defined as

$$\begin{cases} f_u = \inf\{f : |T(f)| \leq 1\} \\ PM = 180^\circ + \angle T(f_u) \\ T = A\beta \end{cases} \quad (9)$$

For the open-loop report characterization, the unity-feedback interpretation sets $\beta = 1$ at the unity-magnitude crossing. The crossing is the first descending $|T(f)| = 1$ crossing, and the phase is evaluated in degrees after phase unwrapping. A non-unity-feedback test must obtain $T = A\beta$ directly. The slew rate is defined as

$$\begin{cases} SR_+ = \max_{t_{in} W} dV_{OUT}/dt \\ SR_- = -\min_{t_{in} W} dV_{OUT}/dt \\ SR = \min(SR_+, SR_-). \end{cases} \quad (10)$$

The common-mode rejection ratio is defined as

$$\begin{cases} A_{cm}(f) = |V_{OUT}/V_{CM}| \\ CMRR(f_{ref}) = 20\log_{10}|A_{vd}(f_{ref})/A_{cm}(f_{ref})| \end{cases} \quad (11)$$

For CMRR, the same AC phasor is applied to VIP and VIN so that the differential input component is zero. For PSRR, a small perturbation is applied to VDD while the input bias is held fixed. Both ratios are evaluated at the row-specific f_{ref} , and the perturbation settings are stored in the row-level record.

The power-supply rejection ratio is defined as

$$\begin{cases} A_{dd}(f) = |V_{OUT}/\Delta V_{DD}| \\ PSRR(f_{ref}) = 20\log_{10}|A_{vd}(f_{ref})/A_{dd}(f_{ref})| \end{cases} \quad (12)$$

The valid common-mode set is

$$\begin{cases} V_{valid} = \{v : E_{out}(v) \leq \varepsilon_{out} \wedge \forall i M_i(v) \geq 0\} \\ ICMR - L = \inf V_{valid} \\ ICMR - H = \sup V_{valid} \end{cases} \quad (13)$$

Where $E_{out}(v)$ denotes the output error relative to the declared reference, and $M_i(v)$ denotes the signed operating-region margin of transistor i . The ICMR endpoints are the infimum and supremum of the continuous valid set. The common-mode sweep step and output-error tolerance are stored in the row-level measurement record.

Total power is given below.

$$P_{total} = \sum_s |V_s I_s| \quad (14)$$

For a single VDD branch, this reduces to $P_{total} = VDD|IDD|$, where VDD is the supply voltage and IDD is the corresponding supply current. Input offset is defined as

$$\begin{cases} VOS(V_{CM}) = V_{ID}|V_{OUT} = V_{REF} \\ VOS(V_{CM}) \approx a + bV_{CM}; L_{VOS} = 1 \end{cases} \quad (15)$$

A row is labeled linear only when its declared R^2 and maximum-residual thresholds are satisfied. For a valid linear fit, the scalar used by the model is the offset at the nominal common-mode voltage of 2.5 V, VOS_{proxy} . The offset remains a separate diagnostic under the conventional op-amp characterization protocol [13].

3.3. Real-cadence data schema and replacement plan

The nominal measurement plan uses one parameterized schematic with the seven grouped variables and fixed L/I_{bias} settings. Each row stores the design variables, analysis configuration, raw-derived metric values, operating-point features, and provenance label. Storing the analysis name, reference frequency, expression version, and convergence flag with each row prevents visually similar curves from being treated as equivalent measurements.

The row-level schema also stores VDD, temperature, load capacitance, process corner, and Monte Carlo seed. For the nominal set, VDD = 5 V, temperature = 27 °C, load capacitance $C_L = 3$ pF, and process corner = TT; the Monte Carlo seed is not applicable to this TT-only evaluation. The internal $C_2 = 3.042$ pF remains independent of $C_L = 3$ pF. PSF identifiers, convergence status, run time, and Calculator/OCEAN expression versions are retained with the source label.

3.4. Reproducible evaluation protocol

The design points are generated and fixed before model training. The 1,400-row grouped working set is split into 1,120 training rows and 280 test rows using the fixed dataset seed, with the report anchor included in the training set. Continuous inputs and outputs are normalized using training-set

statistics only. The direct and physics-guided regressors use identical splits, regularization, and eight-bootstrap procedures, so the comparison is a matched feature ablation.

The nominal calibration protocol is incremental. After confirming the report baseline under the fixed TT/pre-layout configuration, two acquisition arms are evaluated in matched batches: 20 candidates selected by the proposed score and 20 selected by the constrained GP-UCB-style policy. The surrogate models are retrained after each returned batch and evaluated under the same metric expressions, convergence flags, variable schema, and wall-clock accounting. The active-versus-random replay is a separate three-arm comparison with equal batch sizes. This staged acquisition protocol follows the initialization-aware active-learning setting used in [14, 15].

The evaluation procedure consists of the following steps: freeze the grouped-7D split; construct the working rows; train the direct and physics-guided surrogates; run repeated cross-validation; benchmark the six algorithmic entries on the complete LTspice rows; replay active and random selection on the archived pool; and audit provenance. Because the grouped rows are generated within a single bounded neighborhood around one report anchor, this split measures interpolation rather than extrapolation across processes or topologies or an anchor-leave-out condition.

For metric j

$$nMAE_j = MAE_j / s_j \quad (16)$$

Where s_j is the training-fold standard deviation; the macro value is

$$Macro - nMAE = 1/10 \sum_{j=1 \dots 10} nMAE_j. \quad (17)$$

4. Discussion and analysis of results

4.1. Baseline and datasheet provenance




Figure 1. Grouped-7D working-set gain–power feasibility map; green points satisfy the nine scalar specifications (picture credit: original)

The canonical report baseline and the nine scalar specifications plus the VOS-linearity requirement are defined in Section 2.2. Figure 1 summarizes the grouped-7D working set, which contains 956 feasible points, or 68.3%. This ratio is a software and boundary-learning diagnostic; it is not a process-yield estimate.

The grouped-7D working rows remain centered on the report anchor while exposing coupled width and compensation variations. Figure 2 summarizes the corresponding P05–P95 ranges and separates the report anchor from the working distribution. VOS is shown as VOS_{proxy} because the report supplies a linear VOS -versus- VC_{M} trace and a 1.2-mV TT anchor rather than a digitized offset table.

Report Anchor and Grouped-7D Working-Range Summary

(a) Row provenance and (b) metric ranges normalized to the report anchor




Figure 2. Report-anchor and grouped-7D working-range summary. (a) Row provenance; (b) metric ranges normalized to the report anchor, with blue intervals denoting P05–P95 and red diamonds denoting report values (picture credit: original)

4.2. Surrogate prediction performance




Figure 3. Direct versus physics-guided Ridge MAE on the grouped-7D test split; bars show metric-wise MAE (picture credit: original)

On the grouped-7D test split, the physics-guided model reduces macro normalized mean absolute error (MAE) from 0.075 to 0.0664. As shown in Figure 3, the largest relative reductions occur for CMRR, VOS_{proxy} and SR; PM and the two ICMR limits are essentially unchanged. Figure 3 also shows lower displayed errors for Power and the ICMR limits, whereas SR and VOS_{proxy} have larger errors. This supports a metric-dependent benefit from operating-region features, not a universal improvement claim.

4.3. Portable LTspice algorithm baseline

The independent audit uses 32 complete rows from the archived LTspice 17.2.4 data and five repeats of four-fold cross-validation, as shown in Table 2. The remaining 64 rows are partial because they do not contain all ten output labels and are excluded without imputation. All six entries are evaluated under the same complete-row partitions and ten-output feasibility definition. Mean nMAE is the macro-average over the ten outputs after training-fold standardization; area under the precision-recall curve (AUC-PR), ECE, and Brier score evaluate the predicted feasible label.

Table 2. Portable LTspice algorithm-audit baselines; cGP-UCB is a policy/acquisition reference

Model	nMAE	AUC-PR	ECE	Brier
Direct Ridge	0.507	0.946	0.241	0.166
Physics Ridge	0.525	0.936	0.259	0.184
ExtraTrees	0.497	0.948	0.230	0.122
Physics ExtraTrees	0.503	0.942	0.232	0.121
MLP	0.413	0.980	0.219	0.122
cGP-UCB-style	0.393	0.936	0.262	0.174

The constrained GP-UCB-style entry is a policy/acquisition reference rather than a conventional regression model. Its nMAE is reported only as an audit value under the same complete-row split and ten-output scoring protocol; it is not treated as a direct regression comparison. The five predictive baselines and the policy reference therefore serve different diagnostic roles, and the portable results do not establish universal algorithmic superiority for the ST02 grouped-7D circuit.

Within the current bounded working distribution, fold-to-fold variation is low for most metrics. This result describes split stability within that distribution; it does not measure transfer to another process, topology, or measurement regime.

4.4. Feature audit and error interpretation

Figure 4 is a feature audit of the working data, not a device-physics law. The strongest positive associations with gain are observed for $gmro$ ($r = 0.91$), headroom ($r = 0.86$), and the saturation ratio ($r = 0.72$). In the SR column, the strongest association is with the MN5 saturation proxy ($r = 0.71$). The audit contains no direct compensation feature, so it does not support a compensation-specific association or a causal interpretation. Figure 4 provides only the feature-correlation audit; it does not isolate compensation or mismatch effects, and no mismatch descriptors are included in the reported ablation.




Figure 4. Pearson correlation between grouped-7D physics proxies and output metrics; this is a feature-audit diagnostic (picture credit: original)

4.5. Round-1 active-sampling replay

As shown in Figure 5, the first round ranks 5,000 grouped-7D candidates and selects 20 diverse points. The macro normalized MAE is 0.0664 before addition, 0.0664 after the active addition, and 0.0664 after the equal-size random addition. The active batch is therefore interpreted as a boundary-probing queue, not as evidence of a global error reduction. No GP-UCB result is inserted into the replay comparison without a matched completed batch.




Figure 5. Round-1 grouped-7D active-sampling scores for the 20 selected candidates (picture credit: original)

Operationally, the active score is a queueing rule for expensive simulations. At round t , bootstrap models evaluate uncertainty, normalized boundary proximity, predicted feasibility, and novelty on the unlabeled grouped-7D pool; a standardized-distance filter removes redundant candidates; and the top 20 parameter sets are exported with predicted margins. The selected score is useful for

organizing a matched nominal experiment, while the replay comparison is the only result claimed here.

5. Conclusion

This paper presented a provenance-aware, physics-guided feasibility-region learning and active-sampling workflow for a two-stage CMOS operational amplifier. The workflow uses the confirmed seven grouped physical variables, ten output fields including VOS_{proxy} , nine operating-region features, bootstrap uncertainty, algorithm baselines, calibration diagnostics, and equal-budget replay. The controlled study illustrates the importance of explicit measurement definitions and matched validation: the report anchor defines the nominal circuit, the grouped-7D working set checks the learning interface, and the simulator-based reference rows perform an independent algorithm audit.

The evidence is limited to the nominal TT/pre-layout setting. The report provides the real schematic anchor; the grouped-7D rows and VOS_{proxy} are extensions of the working data; and the portable LTspice experiment is a stand-alone algorithm reference. The claims are limited to the working-data interpolation track. RQ2 is evaluated as an explicitly controlled calibration protocol, and the archived replay shows no global MAE gain. The process-specific replacement contract is defined in Section 3. PVT, mismatch, layout parasitics, and yield are outside the scope of the present claim.

References

- [1] Settaluri, K., Haj-Ali, A., Huang, Q., Hakhamaneshi, K., & Nikolic, B. (2020). AutoCkt: Deep reinforcement learning of analog circuit designs. In *Proceedings of the DATE*, Grenoble, France, March 9–13, pp. 490–495.
- [2] Wang, H., Wang, K., Yang, J., Shen, L., Sun, N., Lee, H.-S., & Han, S. (2020). GCN-RL circuit designer: Transferable transistor sizing with graph neural networks and reinforcement learning. In *Proceedings of the DAC*, San Francisco, CA, USA (virtual), July 20–24, pp. 1–6.
- [3] Wu, Z., & Savidis, I. (2020). CALT: Classification with adaptive labeling thresholds for analog circuit sizing. In *Proceedings of the ACM/IEEE MLCAD*, Reykjavík, Iceland (virtual), November 16–20, pp. 49–54.
- [4] Cao, W., Benosman, M., Zhang, X., & Ma, R. (2022). Domain knowledge-infused deep learning for automated analog/RF circuit parameter optimization. In *Proceedings of the DAC*, San Francisco, CA, USA, July 10–14, pp. 1015–1020.
- [5] Choi, Y., Choi, M., Lee, K., & Kang, S. (2023). MA-Opt: Reinforcement learning-based analog circuit optimization using multi-actors. In *Proceedings of the DATE*, Antwerp, Belgium, April 17–19, pp. 1103–1108.
- [6] Gu, T., Wang, J., Bi, Z., Yan, C., Yang, F., Qin, Y., Cui, T., & Zeng, X. (2024). tSS-BO: Scalable Bayesian optimization for analog circuit sizing via truncated subspace sampling. In *Proc. DATE*, Valencia, Spain, March 25–27, pp. 535–540.
- [7] Li, Y., Lin, Y., Madhusudan, M., Sharma, A., Sapatnekar, S., Harjani, R., & Hu, J. (2021). A circuit attention network-based actor-critic learning approach to robust analog transistor sizing. In *Proc. ACM/IEEE MLCAD*, Raleigh, NC, USA (virtual), 30 August–3 September, pp. 101–106.
- [8] Uhlmann, Y., Brunner, M., Bramlage, L., Scheible, J., & Curio, C. (2023). Procedural- and reinforcement-learning-based automation methods for analog integrated circuit sizing in the electrical design space. *Electronics*, 12(2), Art. no. 302.
- [9] Chen, C., Wang, H., Song, X., Liang, F., Wu, K., & Tao, T. (2022). High-dimensional Bayesian optimization for analog integrated circuit sizing based on dropout and gm/ID methodology. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(11), 4808–4820.
- [10] Touloupas, K., & Sotiriadis, P. P. (2022). Mixed-variable Bayesian optimization for analog circuit sizing through device representation learning. *Electronics*, 11, Art. no. 3127.
- [11] Abuelnasr, A., Ragab, A., Amer, M., Gosselin, B., & Savaria, Y. (2024). Incremental reinforcement learning for multi-objective analog circuit design acceleration. *Engineering Applications of Artificial Intelligence*, 129, Art. no. 107426.

- [12] Lyu, R., Zhao, A., Meng, Y., Zhu, K., Bi, Z., Yan, C., Yang, F., Zhou, D., & Zeng, X. (2024). Revisiting sensitivity-based analog sizing with derivative-aware Bayesian optimization and error-suppressed adjoint analysis. In *Proc. ACM/IEEE ICCAD*, Newark, NJ, USA, October 27-31, Vol. 138, pp. 1-9.
- [13] Aueamnuay, C., Kayyil, A. V., Liu, J., Thota, N. B., & Allstot, D. J. (2020). gm/ID design considerations for subthreshold-based CMOS two-stage operational amplifiers. In *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Seville, Spain (virtual), 10–21 October 2020, pp. 1–5.
- [14] Kircali Ata, S., Kong, Z.-H., James, A., Cai, L., Yeo, K. S., Aung, K. M. M., Foo, C. S., & James, A. (2024). The initialization factor: Understanding its impact on active learning for analog circuit design. In *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Singapore, 1May 19-22, pp. 1–5.
- [15] Kircali Ata, S., Kong, Z.-H., James, A., Cai, L., Yeo, K. S., Aung, K. M. M., Foo, C. S., & James, A. (2024). Towards safe and efficient analog circuit design: Active learning for feasibility region exploration. In *Proc. IEEE Asia Pacific Conf. Circuits Syst. (APCCAS)*, Taipei, Taiwan, November 7-9, pp. 193–197.