

Machine Learning-Based Estimation of Cumulants of Chiral Condensate via Multi-Ensemble Reweighting with Deborah.jl



Benjamin J. Choi ^{a,*}, Hiroshi Ohno ^a and Akio Tomiya ^{b,c,d}

^aCenter for Computational Sciences, University of Tsukuba,
1-1-1 Tennodai, Tsukuba, Ibaraki 305-8577, Japan

^bDepartment of Information and Mathematical Sciences, Tokyo Woman's Christian University,
2-6-1 Zempukuji, Suginami-ku, Tokyo 167-8585, Japan

^cRIKEN Center for Computational Science,
7-1-26 Minatojima-minami-machi, Chuo-ku, Kobe 650-0047, Japan

^dDepartment of Physics, Kyoto University, Kitashirakawa, Sakyo-ku, Kyoto 606-8502, Japan
E-mail: benchoi@ccs.tsukuba.ac.jp, hohno@ccs.tsukuba.ac.jp,
akio@yukawa.kyoto-u.ac.jp

We investigate a bias-corrected machine learning (ML) strategy for estimating traces of the inverse Dirac operator, $\text{Tr } M^{-n}$ ($n = 1, 2, 3, 4$), motivated by the need for higher-order cumulants of the chiral condensate near the finite-temperature QCD critical endpoint. Our supervised regression framework is trained on Wilson-clover ensembles with the Iwasaki gauge action, and we explore two input feature scenarios: one using $\text{Tr } M^{-1}$ and another relying solely on gauge observables (plaquette and rectangle), enabling a fully feature-based prediction pipeline. Using $\text{Tr } M^{-1}$ both as a physical input to cumulant construction and as a feature for predicting higher powers, we find that even with $\sim 1\%$ labeled data, the resulting susceptibility, skewness, and kurtosis remain statistically consistent with fully measured baselines, reducing computational cost to about 26%. In the feature-only approach, where correlations rather than explicit stochastic traces drive the predictions, bias correction plays a more pronounced role. We quantify this impact through multi ensemble reweighting across nearby quark masses. Our results demonstrate that bias-corrected ML estimates can significantly reduce measurement overhead while preserving the stability of higher-order observables relevant for locating the QCD critical endpoint. Code for this work is available at   Deborah.jl.

The 42nd International Symposium on Lattice Field Theory (LATTICE2025)
2-8 November 2025
Tata Institute of Fundamental Research, Mumbai, India

*Speaker

1. Introduction



Understanding the critical endpoint in the finite-temperature QCD phase diagram requires the precise determination of higher-order fluctuations of the chiral condensate, which acts as the order parameter [1, 2]. Techniques such as the kurtosis intersection method rely on these fluctuations to identify the critical behavior and extract universal scaling properties near the transition [3, 4].

The evaluation of these fluctuations involves computing traces of powers of the inverse Dirac operator, $\text{Tr } M^{-n}$. Even when stochastic estimators such as Hutchinson-based approaches are used [5], the dominant computational effort comes from repeated solutions of large sparse linear systems with iterative solvers. As a result, brute-force measurement at high statistics remains costly.

Recent developments in artificial intelligence for scientific applications have opened new possibilities for accelerating computationally intensive studies in lattice field theory [6]. In particular, machine-learning-assisted approaches have begun to be explored as a means to reduce the cost of stochastic measurements while maintaining physics fidelity.

A recent strategy to mitigate this computational burden applies supervised machine learning together with a bias correction scheme motivated by the All Mode Averaging (AMA) framework [7, 8], as demonstrated in Ref. [9]. In this approach, explicitly measured configurations form a labeled set that is divided into training and bias-correction subsets, while the remaining unlabeled configurations are used for prediction, balancing model accuracy against bias control.

In this work, we systematically vary both the size of the labeled dataset and the training fraction within it, assessing how these choices influence the precision of ML-based estimators for $\text{Tr } M^{-n}$. Beyond the direct trace estimates, we explore a physics application of the ML outputs through multi-ensemble reweighting across nearby quark masses. These tests probe whether the ML-derived observables retain sufficient fidelity for thermodynamic analyses relevant to the QCD critical endpoint.

We further compare bias-corrected and uncorrected setups to isolate the impact of bias removal. Conventional full-statistics evaluations are carried out in parallel to provide a direct baseline for judging the accuracy and stability of the ML-assisted estimation workflow. The implementation used in this study is provided within the   Deborah.jl framework [10].

2. Formalism, Dataset, and ML Estimation Framework

2.1 Notation, Dataset, and Cumulant Observables

For convenience, Table 1 provides the notation and definitions that will be used throughout this paper. These symbols will be referenced repeatedly in later sections, and the table serves as a compact glossary rather than repeating explanations inline.

To test the machine learning (ML) estimation strategy in a practical lattice-QCD environment, we utilize gauge configurations originally produced for Ref. [11]. The ensembles were generated on the Oakforest-PACS supercomputing system [12] using the BQCD framework [13], with $N_f = 4$ Wilson-Clover fermions [14] and the Iwasaki gauge action [15, 16]. This dataset serves as a well understood reference environment for evaluating ML-based trace estimation performance. The specific simulation parameters relevant to the present analysis are listed in Table 2.

Symbol	Description
X	input observables (or feature) (e.g., $X = \text{Plaquette, Rectangle, Tr } M^{-1}$)
Y	output observables (or target) (e.g., $Y = \text{Tr } M^{-n}$ for $n = 1, 2, 3, 4$)
S^Z	the total dataset of $Z = X, Y$ where $S^Z = S_{\text{LB}}^Z \cup S_{\text{UL}}^Z$
S_{LB}^Z	the labeled set of the original data for $Z = X, Y$ where $S_{\text{LB}}^Z = S_{\text{TR}}^Z \cup S_{\text{BC}}^Z$
S_{TR}^Z	the training set of the original data for $Z = X, Y$
S_{BC}^Z	the bias correction set of the original data for $Z = X, Y$
S_{UL}^Z	the unlabeled set of the original data for $Z = X, Y$
N	the number of elements of S^Z where $N = S^X = S^Y $
N_{LB}	the number of elements of S_{LB}^Z where $N_{\text{LB}} = S_{\text{LB}}^X = S_{\text{LB}}^Y $
N_{TR}	the number of elements of S_{TR}^Z where $N_{\text{TR}} = S_{\text{TR}}^X = S_{\text{TR}}^Y $
N_{BC}	the number of elements of S_{BC}^Z where $N_{\text{BC}} = S_{\text{BC}}^X = S_{\text{BC}}^Y $
N_{UL}	the number of elements of S_{UL}^Z where $N_{\text{UL}} = S_{\text{UL}}^X = S_{\text{UL}}^Y $
$f(X)$	the model trained with S_{TR}^X and S_{TR}^Y
Y^P	the ML estimation on Y
S_{BC}^P	the bias correction set composed of the ML estimations Y_{BC}^P corresponding to S_{BC}^Y
S_{UL}^P	the unlabeled set composed of the ML estimations Y_{UL}^P corresponding to S_{UL}^Y

Table 1: Notation and convention used in this paper for the explanation of our work.

ID	$N_S^3 \times N_T$	β	c_{sw}	κ	N
L12T4b1.60k13575	$12^3 \times 4$	1.60	2.065	0.13575	20000
L12T4b1.60k13577	$12^3 \times 4$	1.60	2.065	0.13577	20000
L12T4b1.60k13580	$12^3 \times 4$	1.60	2.065	0.13580	20000
L12T4b1.60k13582	$12^3 \times 4$	1.60	2.065	0.13582	20000
L12T4b1.60k13585	$12^3 \times 4$	1.60	2.065	0.13585	20000

Table 2: Data used in this paper, originally produced for Ref. [11]. Here, N follows the same convention as in Table 1, where it corresponds to the number of gauge configurations.

We evaluate the impact of ML-based trace predictions on thermodynamic observables derived from the Wilson-Clover operator,

$$M(x, y) = \frac{1}{2K} \delta_{x,y} + \frac{i}{4} c_{\text{sw}} \sigma_{\mu\nu} F_{\mu\nu}(x) \delta_{x,y} - \frac{1}{2} \sum_{\mu=1}^4 \sum_{s=\pm 1} (1 - s \gamma_{\mu}) U_{s\mu}(x) \delta_{x, y+s\hat{\mu}}. \quad (1)$$

From the traces $\text{Tr } M^{-n}$, we construct the quark-loop operators,

$$\begin{aligned} Q_1 &= N_f \text{Tr } M^{-1}, & Q_2 &= -N_f \text{Tr } M^{-2} + (N_f \text{Tr } M^{-1})^2, \\ Q_3 &= 2N_f \text{Tr } M^{-3} - 3(N_f \text{Tr } M^{-2})(N_f \text{Tr } M^{-1}) + (N_f \text{Tr } M^{-1})^3, \\ Q_4 &= -6N_f \text{Tr } M^{-4} + 8(N_f \text{Tr } M^{-3})(N_f \text{Tr } M^{-1}) + 3(N_f \text{Tr } M^{-2})^2 \\ &\quad - 6(N_f \text{Tr } M^{-2})(N_f \text{Tr } M^{-1})^2 + (N_f \text{Tr } M^{-1})^4. \end{aligned} \quad (2)$$

Using ensemble averages of Q_i , we obtain the cumulants of the chiral condensate,

$$\Sigma = C_1/V, \quad \chi = C_2/V, \quad S = C_3/C_2^{3/2}, \quad K = C_4/C_2^2, \quad (3)$$

with

$$\begin{aligned} C_1 &= \langle Q_1 \rangle, \quad C_2 = \langle Q_2 \rangle - \langle Q_1 \rangle^2, \quad C_3 = \langle Q_3 \rangle - 3\langle Q_2 \rangle \langle Q_1 \rangle + 2\langle Q_1 \rangle^3, \\ C_4 &= \langle Q_4 \rangle - 4\langle Q_3 \rangle \langle Q_1 \rangle - 3\langle Q_2 \rangle^2 + 12\langle Q_2 \rangle \langle Q_1 \rangle^2 - 6\langle Q_1 \rangle^4. \end{aligned} \quad (4)$$

2.2 Supervised Learning with Bias Correction

In our supervised ML setup, each configuration carries input features X and target observables Y . Only the labeled subset $S_{\text{LB}}^Z \subset S^Z$ ($Z = X, Y$) has both X and Y available. We split this labeled set into a training part S_{TR}^Z and a bias-correction part S_{BC}^Z , while the remaining configurations form the unlabeled set S_{UL}^Z . The basic two-step procedure is

1. Train a regression model $f(X)$ using S_{TR}^X and S_{TR}^Y .
2. Apply f to $X \in S_{\text{UL}}^X$ to obtain predictions $Y^P = f(X) \approx Y$.

Because the model is trained on only a subset of the full data, it can develop systematic biases. Following the bias-correction strategy of Ref. [9], we use S_{BC}^Z to correct this bias and construct the estimator

$$\bar{Y}_{\mathcal{P}1} = \frac{1}{N_{\text{UL}}} \sum_{Y_i \in S_{\text{UL}}^Y} Y_i^P + \frac{1}{N_{\text{BC}}} \sum_{Y_j \in S_{\text{BC}}^Y} (Y_j - Y_j^P), \quad (5)$$

where N_{UL} and N_{BC} denote the sizes of S_{UL}^Y and S_{BC}^Y , respectively. A second estimator, $\bar{Y}_{\mathcal{P}2}$, was also explored in Ref. [17] but is not considered here; see that work for its definition.


2.3 Scanning the Labeled and Training Fractions

To control the amount of exact CG work and the size of the bias-correction sample, we parametrize the partition in terms of the fractions

$$\mathcal{R}_{\text{LB}} \equiv \frac{N_{\text{LB}}}{N}, \quad \mathcal{R}_{\text{TR}} \equiv \frac{N_{\text{TR}}}{N_{\text{LB}}}, \quad (6)$$

with $N = N_{\text{LB}} + N_{\text{UL}}$ and $N_{\text{LB}} = N_{\text{TR}} + N_{\text{BC}}$. We explore a set of representative values, $\mathcal{R}_{\text{LB}} \in \{1, 2, \dots, 24, 25\%\}$ and $\mathcal{R}_{\text{TR}} \in \{0, 10, \dots, 90, 100\%\}$, in order to map out how the ML estimation quality depends on the amount of labeled data and on the training fraction.

The extreme choices $\mathcal{R}_{\text{TR}} = 0\%$ and $\mathcal{R}_{\text{TR}} = 100\%$ provide useful benchmarks. In the former case, all labeled configurations are reserved for a purely ‘‘conventional’’ estimate without ML, allowing us to monitor convergence as \mathcal{R}_{LB} increases. In the latter case, all labeled data are used for training and no bias correction is applied, so the impact of omitting the correction can be directly assessed.

(a) $\kappa = 0.13575$, L12T4b1.60k13575 (the heaviest quark) (b) $\kappa = 0.13590$, L12T4b1.60k13585 (the lightest quark)**Figure 1:** Correlation between physical observables.

2.4 Correlation Structure among Observables

Supervised regression benefits from strong correlations between input features and target observables. In Fig. 1 we show the correlation matrix among $\text{Tr} M^{-n}$, the plaquette, the rectangle, and the Polyakov loop for two representative ensembles, L12T4b1.60k13575 (heaviest quark) and L12T4b1.60k13585 (lightest quark). Apart from $\text{Tr} M^{-4}$, which exhibits weaker correlations and thus requires some care, most pairs display substantial positive correlation, which motivates the use of these quantities as ML inputs and targets.

2.5 Two Practical ML Setups

The traces $\text{Tr} M^{-n}$ entering cumulants are evaluated by a Hutchinson-type stochastic trace estimator, in which random noise vectors probe the operator, and the dominant cost arises from CG inversions of the Dirac matrix. To further reduce this cost, we consider two complementary choices of input features:

\mathcal{F}_{in} In this setup, the feature set \mathcal{F} consists of *internal* trace-based observables, in particular $\text{Tr} M^{-1}$. We use the original CG measurements of $\text{Tr} M^{-1}$ both as a direct input to the cumulant construction and as a feature to predict $\text{Tr} M^{-n}$ ($n = 2, 3, 4$). The ML-predicted higher powers are then combined with the exact $\text{Tr} M^{-1}$ to build the cumulants. This configuration forms the main line of analysis in this work. Because $\text{Tr} M^{-1}$ typically dominates the cumulant expressions, the resulting cumulants remain close to the original CG values even when \mathcal{R}_{LB} is small; at the same time, the required CG work cannot be reduced below the cost of measuring $\text{Tr} M^{-1}$. For cumulants up to kurtosis, this implies a lower bound of about 25% of the baseline cost.

\mathcal{F}_{ex} Here, the feature set \mathcal{F} is built from *external* observables recorded during the HMC evolution, namely the plaquette and rectangle, which enter the Iwasaki gauge action [15, 16] and are available from the original simulations [11]. We use these two quantities as default input features and treat all $\text{Tr} M^{-n}$ as ML targets, so that the quality of the cumulants fully reflects

the ML estimation performance. Although the Polyakov loop is also measured, its correlation with $\text{Tr } M^{-n}$ is weaker than that of the plaquette and rectangle (Fig. 1), and we therefore omit it from the default feature set.

The \mathcal{F}_{in} setup is closer to a ‘‘partial replacement’’ strategy, where only higher powers of M^{-1} are delegated to ML, while the \mathcal{F}_{ex} setup aims at a fully feature-only prediction pipeline. Comparing the two allows us to distinguish the impact of having $\text{Tr } M^{-1}$ explicitly available from the purely feature-based scenario.

2.6 Treatment of Statistical Errors

The configurations analyzed here are generated near a first-order phase transition and exhibit autocorrelations; naive delete-1 jackknife or *i.i.d.* bootstrap procedures that neglect this dependence tend to underestimate statistical errors. A more appropriate treatment uses either delete- g jackknife or block bootstrap resampling.

In our case, the ML estimator combines contributions from subsets of different sizes (unlabeled and bias-correction sets); see Eq. (5). Constructing synchronized jackknife replicas across all these subsets is cumbersome, so we adopt a block bootstrap approach to account for temporal correlations in a straightforward manner. Related discussions of block bootstrap methods in lattice QCD can be found in Ref. [18].

2.7 Evaluation Criterion

To quantify the agreement between the ML-based estimates and the reference results obtained from the conventional CG method, we employ a single scalar metric: the Bhattacharyya coefficient [9]. For two Gaussian distributions, it is defined as

$$C_B(x, r) = \sqrt{\frac{2r}{1+r^2}} \exp\left[-\frac{x^2}{4(1+r^2)}\right], \quad \text{where } x = \frac{|\bar{Y}_{\text{Orig}} - \bar{Y}_{\text{ML}}|}{\sigma_{\text{Orig}}}, \quad r = \frac{\sigma_{\text{ML}}}{\sigma_{\text{Orig}}}. \quad (7)$$

Here, x measures the normalized shift in the mean, and r represents the relative size of the statistical uncertainty of the ML estimation compared with the original result.

The coefficient C_B takes values between 0 and 1, with larger values corresponding to a stronger overlap between the two distributions. In the limiting cases where the ML estimator reproduces both the mean ($x = 0$) and uncertainty ($r = 1$) of the reference result, $C_B = 1$, indicating perfect statistical consistency. Based on typical Gaussian overlap behavior, a value of $C_B \gtrsim 0.95$ corresponds to a separation well below one standard deviation and a relative uncertainty close to unity, and is therefore interpreted here as indicating substantial agreement between the ML prediction and the original measurement. Throughout this work, C_B serves as the primary figure of merit, while x and r are referenced only when a more detailed diagnostic is required.

3. Cumulant estimation with multi-ensemble reweighting

Having established the ML setup and the feature choices, we now proceed to apply the estimators to predict $\text{Tr } M^{-n}$ on all ensembles. These predictions are then combined through multi-ensemble reweighting to extract the cumulants along the quark-mass trajectory. In this setting, the trace

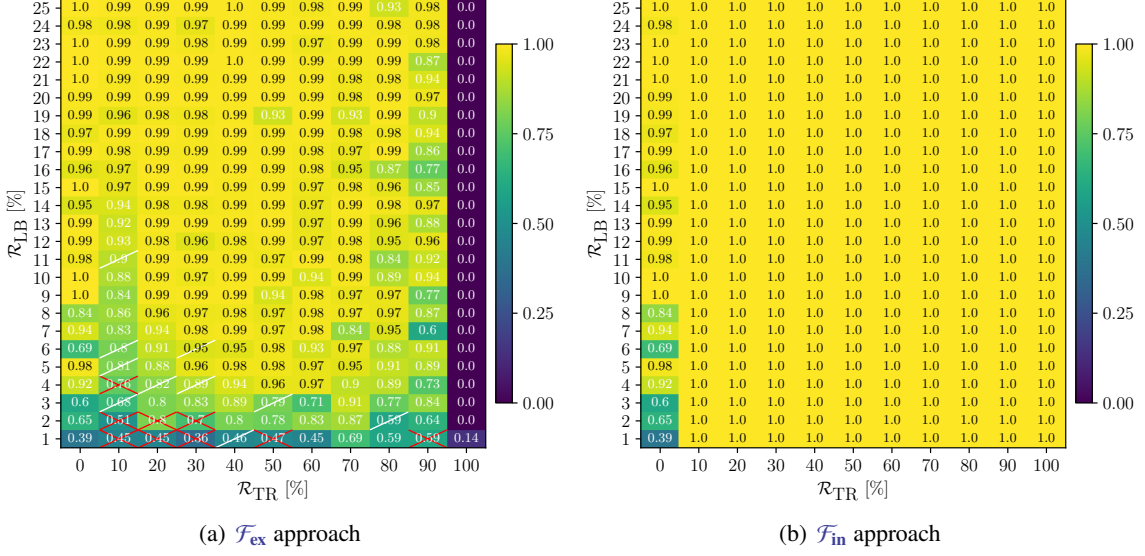


Figure 2: Bhattacharyya coefficient C_B maps for the kurtosis at the transition point $K(\kappa)$ obtained from multi-ensemble reweighting. Panel (a) shows the results for $\mathcal{P}1$ under the \mathcal{F}_{ex} setup. Red cross marks indicate cells where the Newton–Raphson solver failed to converge within the maximum iteration count, while white diagonal marks indicate cases where convergence was achieved but required more than 10 Newton iterations. Panel (b) shows the corresponding $\mathcal{P}1$ results for the \mathcal{F}_{in} setup.

estimates from all available ensembles—generated at common (V, β) but distinct values of κ —are combined to interpolate observables across the quark-mass axis following the standard Ferrenberg–Swendsen reweighting framework [19, 20]. The underlying gauge ensembles were originally produced for Ref. [11] and have previously been used to investigate the finite-temperature phase structure of $N_f = 4$ QCD.

Our implementation combines the usual multi-ensemble machinery with the bias-corrected ML estimators of Eq. (5). For each observable Y (here Y represents the trace combinations entering Q_j), the samples on every ensemble are partitioned into S_{TR}^Y , S_{BC}^Y , S_{UL}^Y and their predicted counterparts S_{BC}^P , S_{UL}^P . We then construct, for each ensemble, four index-aligned sets

$$\begin{aligned} \mathcal{S}_1 &= S_{\text{TR}}^Y \cup S_{\text{BC}}^Y \cup S_{\text{UL}}^Y \equiv \mathcal{S}^Y, & \mathcal{S}_2 &= S_{\text{TR}}^Y \cup S_{\text{BC}}^Y \cup S_{\text{UL}}^P, \\ \mathcal{S}_3 &= S_{\text{TR}}^Y \cup S_{\text{BC}}^Y \equiv \mathcal{S}_{\text{LB}}^Y, & \mathcal{S}_4 &= S_{\text{TR}}^Y \cup S_{\text{BC}}^P, \end{aligned} \quad (8)$$

ensuring that the configuration indices and ordering coincide with the original data. Afterward, the corresponding sets from all κ values are concatenated (*e.g.*, in increasing κ) to form cross-ensemble streams used in the subsequent reweighting steps.

We now discuss the multi-ensemble reweighting results for the \mathcal{F}_{ex} approach. Before presenting the comparable cumulant outputs, we briefly note that the Newton-Raphson solver used to determine the free-energy offsets does not perform uniformly across the scanned $(\mathcal{R}_{\text{LB}}, \mathcal{R}_{\text{TR}})$ parameter space. In particular, a small subset of cases with very limited labeled data ($\mathcal{R}_{\text{LB}} = 1\text{--}4\%$ combined with low \mathcal{R}_{TR}) failed to converge within the allowed iteration limit, while some additional points did converge but only after requiring more than 10 iterations. These behaviors were tracked in advance because they signal instability of the ML-derived input traces rather than an issue with the reweighting procedure itself, and the affected entries are explicitly marked in the following figures.

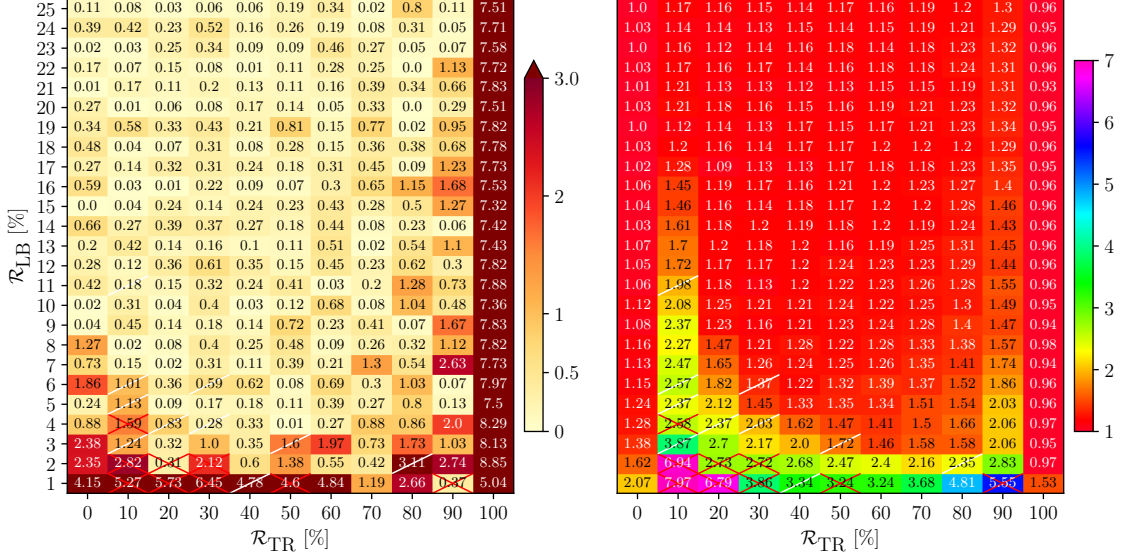


Figure 3: Heatmap showing the $\mathcal{P}1$ results of x and r evaluations, where x and r are defined in Eq. (7), for the estimation of the kurtosis at the phase transition point, $K(\kappa_t)$, based on the \mathcal{F}_{ex} approach. Within the panel, the left half represents the x results, while the right half corresponds to r . In the x maps, lighter shades indicate smaller normalized mean separations x , corresponding to better agreement between the original and predicted means, whereas darker shades represent larger x values, signaling poorer overlap. The kurtosis is computed using $\text{Tr } M^{-n}$ predicted from plaquette and rectangle as the input feature.

With these caveats in place, Fig. 2(a) presents the Bhattacharyya coefficient C_B for the kurtosis estimation at the transition point $K(\kappa_t)$, obtained under the \mathcal{F}_{ex} setup using all ensembles in Table 2.

A clear pattern emerges. First, in the low- \mathcal{R}_{LB} , low- \mathcal{R}_{TR} corner the Newton solver is unreliable and the resulting C_B is also poor. Even nearby cells without explicit convergence issues tend to show suppressed overlap, reflecting the limited quality of S_{UL}^P when the labeled fraction is too small. For this dataset the problematic region largely disappears once $\mathcal{R}_{\text{LB}} \gtrsim 20\%$. Second, the combination of small \mathcal{R}_{LB} and large \mathcal{R}_{TR} (*i.e.* very small \mathcal{R}_{BC}) again leads to reduced C_B . In particular, the column $\mathcal{R}_{\text{TR}} = 90\%$ only reaches stable high-overlap values once $\mathcal{R}_{\text{LB}} \gtrsim 23\%$, similar to the behavior at $\mathcal{R}_{\text{TR}} = 10\%$. This indicates that the ML quality degrades both when the training set is extremely small and when the bias-correction set becomes too small.

The most striking feature appears in the $\mathcal{R}_{\text{TR}} = 100\%$ column, where the labeled data are used exclusively for training and no bias correction is performed. Here C_B is essentially zero throughout the scan. The x (see Eq. (7)) analysis in Fig. 3 shows that the normalized mean separation reaches $x \approx 7$ – 8 , indicating that the ML result for $K(\kappa_t)$ deviates by several standard deviations from the conventional reweighting result. Similar behavior, though with varying magnitude, is seen in other datasets (not shown), systematically confirming that removing the bias correction at the reweighting stage leads to significant distortions once higher-order cumulants and the transition point are inferred.

For comparison, Fig. 2(b) displays the corresponding multi-ensemble reweighting results for the \mathcal{F}_{in} approach. In this case the $\mathcal{P}1$ estimator yields $C_B \approx 1$ essentially everywhere in the scanned $(\mathcal{R}_{\text{LB}}, \mathcal{R}_{\text{TR}})$ range, except for the low- \mathcal{R}_{LB} region at $\mathcal{R}_{\text{TR}} = 0\%$, where only the labeled set is used and the statistical precision is inevitably worse than for the full original data. Thus, when the exact

$\text{Tr } M^{-1}$ is retained as in the \mathcal{F}_{in} setup, the reweighting-based cumulants and the extracted transition point remain remarkably robust against changes in $(\mathcal{R}_{\text{LB}}, \mathcal{R}_{\text{TR}})$, whereas in the fully ML-driven \mathcal{F}_{ex} case, bias correction plays a crucial role in stabilizing the multi-ensemble cumulant analysis.

4. Conclusion

In this work we have applied the bias-corrected ML framework of Ref. [9] to the estimation of thermodynamic observables relevant to finite-temperature QCD, and compared the results with conventional calculations performed on the full set of stochastic trace measurements.

For the \mathcal{F}_{in} approach, the ML-assisted estimates reproduce the original results for traces, single-ensemble cumulants, and multi-ensemble reweighting almost perfectly across all tested datasets and parameter pairs, including the most extreme case with $\mathcal{R}_{\text{LB}} = 1\%$. This robustness is naturally understood from the fact that the full set of measured $\text{Tr } M^{-1}$ values is kept, and the cumulants are dominated by this observable. Under this assumption, the total cost can, in our setup, be reduced to about

$$\frac{100 + 1 + 1 + 1}{400} \simeq 25.75\%$$

of the original measurement budget while retaining the same precision, suggesting a sizable practical gain if the observed stability persists in broader applications.

The \mathcal{F}_{ex} approach is more ambitious: all $\text{Tr } M^{-n}$'s are inferred from gauge observables instead of lower-order traces. Consequently, the overlap with the full-data results, quantified by the Bhattacharyya coefficient C_B , shows a stronger dependence on the size of the labeled set. In particular, for the ensembles studied here, C_B improves systematically with increasing \mathcal{R}_{LB} and becomes more stable once $\mathcal{R}_{\text{LB}} \gtrsim 20\%$, indicating a potential cost reduction to the $\sim 20\%$ level. At the same time, the use of independent gauge features makes this setup more sensitive to modeling choices, and careful validation over additional ensembles and possibly more sophisticated ML architectures will be needed before drawing firm conclusions.

A lesson from the \mathcal{F}_{ex} analysis is the importance of bias correction, especially when ML outputs are fed into multi-stage workflows. In the extreme case $\mathcal{R}_{\text{TR}} = 100\%$, where the labeled data are used only for training and no bias correction is performed, the agreement with the full-data results deteriorates markedly as one proceeds from trace estimation to multi-ensemble reweighting and finally to the determination of the kurtosis at the transition point κ_t . The accumulated shift in this “no-bias-correction” column reaches many standard deviations, underscoring that even small residual biases in intermediate ML predictions can be amplified through higher-order observables and interpolation procedures.

In summary, the \mathcal{F}_{in} strategy appears to be a practical and robust option for reducing the computational footprint of fermionic observables, at least for the ensembles and observables considered in this study. The \mathcal{F}_{ex} strategy is more challenging but also potentially more flexible, and our results indicate that its reliable use will require both a sufficient labeled fraction and explicit bias correction at the cumulant level.

We expect that further tests on different lattices and actions, and extensions to other thermodynamic and fluctuation observables, will clarify how broadly this ML-based framework can be deployed within the lattice QCD community.

Acknowledgments

B. J. C. would like to thank Takayuki Sumimoto for his early contributions and dedication to the initial stage of this work. He also thanks Ho Hsiao for fruitful discussions. The work of A. T. was partially supported by JSPS KAKENHI Grants No. 20K14479, No. 22H05111, No. 22K03539 and JST BOOST, Japan Grant No. JPMJBY24F1. A. T. and H. O. were partially supported by JSPS KAKENHI Grant No. 22H05112. B. J. C. and part of this work were supported by MEXT as “Program for Promoting Researches on the Supercomputer Fugaku” (Grant Number JPMXP1020230411, JPMXP1020230409).

References

- [1] O. Philipsen *Symmetry* **13** (2021) 2079 [2111.03590].
- [2] J.N. Guenther *Eur. Phys. J. A* **57** (2021) 136 [2010.15503].
- [3] X.-Y. Jin *et al.* *Phys. Rev. D* **91** (2015) 014508 [1411.7461].
- [4] Y. Kuramashi *et al.* *Phys. Rev. D* **94** (2016) 114507 [1605.04659].
- [5] S.-J. Dong and K.-F. Liu *Phys. Lett. B* **328** (1994) 130 [hep-lat/9308015].
- [6] A. Tomiya *J. Phys. Soc. Jap.* **94** (2025) 031006.
- [7] G.S. Bali, S. Collins and A. Schafer *Comput. Phys. Commun.* **181** (2010) 1570 [0910.3970].
- [8] T. Blum, T. Izubuchi and E. Shintani *Phys. Rev. D* **88** (2013) 094503 [1208.4349].
- [9] B. Yoon, T. Bhattacharya and R. Gupta *Phys. Rev. D* **100** (2019) 014504 [1807.05971].
- [10] B.J. Choi, *Deborah.jl*, Feb., 2026. 10.5281/zenodo.18755146.
- [11] H. Ohno *et al.* *PoS LAT2018* (2018) 174 [1812.01318].
- [12] T. Boku, K.-I. Ishikawa, Y. Kuramashi and L. Meadows, 1709.08785.
- [13] Y. Nakamura and H. Stuben *PoS LAT2010* (2010) 040 [1011.0199].
- [14] B. Sheikholeslami and R. Wohlert *Nucl. Phys. B* **259** (1985) 572.
- [15] Y. Iwasaki *Nucl. Phys. B* **258** (1985) 141.
- [16] Y. Iwasaki, 1111.7054.
- [17] B.J. Choi *et al.* *PoS LATTICE2024* (2024) 033 [2411.18170].
- [18] RBC, UKQCD collaboration *Phys. Rev. D* **111** (2025) 074514 [2409.11379].
- [19] A.M. Ferrenberg and R.H. Swendsen *Phys. Rev. Lett.* **61** (1988) 2635.
- [20] A.M. Ferrenberg and R.H. Swendsen *Phys. Rev. Lett.* **63** (1989) 1195.