Boltzmann sampling with quantum annealers via fast Stein correction

Ryosuke Shibukawa,[1] Ryo Tamura,[1,2,3,*] and Koji Tsuda[1,2,3,†]

[1]*Graduate School of Frontier Sciences, The University of Tokyo, Chiba 277-8568, Japan*
[2]*Center for Basic Research on Materials, National Institute for Materials Science, Ibaraki 305-0044, Japan*
[3]*RIKEN Center for Advanced Intelligence Project, Tokyo 103-0027, Japan*

(Received 8 September 2023; accepted 30 September 2024; published 21 October 2024)

Despite the attempts to apply a quantum annealer to Boltzmann sampling, it is still impossible to perform accurate sampling at arbitrary temperatures. Conventional distribution correction methods such as importance sampling and resampling cannot be applied, because the analytical expression of sampling distribution is unknown for a quantum annealer. Stein correction [Q. Liu and J. Lee, in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research (PMLR, 2017), Vol. 54, pp. 952–961] can correct the samples by weighting without the knowledge of the sampling distribution, but the naive implementation requires the solution of a large-scale quadratic program, hampering usage in practical problems. In this article, a fast and approximate method based on a random feature map and exponentiated gradient updates is developed to compute the sample weights and is used to correct the samples generated by D-Wave quantum annealers. In benchmarking problems, it is observed that the residual error of thermal average calculations is reduced significantly. If combined with our method, quantum annealers may emerge as a viable alternative to long-established Markov chain Monte Carlo methods.

DOI: 10.1103/PhysRevResearch.6.043050

I. INTRODUCTION

Boltzmann sampling of the Ising model is central in the studies of critical phenomena [1,2] and machine learning [3,4]. Although Markov chain Monte Carlo (MCMC) methods can generate samples according to Boltzmann distribution, they often fall short for large models due to slow mixing [5]. To efficiently perform Boltzmann sampling, various improvement techniques such as exchange Monte Carlo and population annealing have been proposed in statistical physics [6–8].

As an alternative, quantum annealers (QAs) [9] have been expected to work as a means to achieve accurate Boltzmann sampling [10–15]. Theoretically, it has been shown that the distribution of quantum annealing samples deviates from Boltzmann distribution [16]. Nevertheless, scientific discussion is still unsettled about whether QA samples can be used as samples of a Boltzmann distribution in practical terms. Recently, Nelson *et al.* argued that the D-Wave quantum annealer works as an accurate sampler at certain temperatures [10], but it does not work well at arbitrary temperatures.

To accurately obtain the target distribution, conventional distribution correction techniques such as importance sampling and resampling [5] are the powerful tools when the original distribution is known [17]. However, the sampling distribution by QAs has not been analytically described, and these conventional approaches cannot be applied to achieve Boltzmann sampling at arbitrary temperatures. To solve this problem, a correction method without knowing the original distribution is needed when only the target distribution (in this case, Boltzmann distribution with arbitrary temperature) is given.

Liu and Lee proposed a "black-box" distribution correction method based on Stein statistics, where the analytical form of the original distribution is not needed [18]. The original samples are assigned the weights to fit to the target distribution via quadratic programming. In the original paper, the theoretical properties are largely unsolved, but Hodgkinson *et al.* showed the convergence of Stein correction for samples generated by a Markov chain [19].

This study investigates how well the Stein correction works in the distribution correction of QA samples to a Boltzmann distribution with a given temperature. First, we develop a fast approximate algorithm of Stein correction, because the $O(n^3)$ computational cost of quadratic programming for the number of samples, n, is prohibitive for a large number of samples. Next, in benchmarking studies, we observe that the estimation error of internal energy, magnetic susceptibility, and Binder cumulant of some Ising models decreases in a large extent by Stein correction. This result implies that Stein correction is useful for improving the sample quality for applications such as critical phenomena and machine learning.

II. METHOD

A. Boltzmann sampling

We are engaged in sampling from the Boltzmann distribution $p(\boldsymbol{x}) \sim \exp[-\beta H_{\mathrm{Ising}}(\boldsymbol{x})]$, $\boldsymbol{x} \in \{-1, 1\}^d$, where the

*Contact author: tamura.ryo@nims.go.jp
†Contact author: tsuda@k.u-tokyo.ac.jp

Hamiltonian is described as

$$H_{\mathrm{Ising}}(\boldsymbol{x}) = -\sum_{i,j\in E} J_{ij}x_i x_j - \sum_{i\in V} h_i x_i. \tag{1}$$

β denotes the inverse temperature, $V \subseteq [1, d]$, and $E \subseteq [1, d] \times [1, d]$. Here, we assume that the parameters J_{ij} and h_i are in the range of -1 and 1. The thermal average of observable $\mathcal{O}(\boldsymbol{x})$ is defined by

$$\langle \mathcal{O}(\boldsymbol{x})\rangle_\beta = \frac{\mathrm{Tr}\mathcal{O}(\boldsymbol{x})\exp[-\beta H_{\mathrm{Ising}}(\boldsymbol{x})]}{\mathrm{Tr}\exp[-\beta H_{\mathrm{Ising}}(\boldsymbol{x})]}. \tag{2}$$

Since the trace calculation is impossible for large models, this trace is approximately replaced to the average of some samples. When $\beta H_{\mathrm{Ising}}(\boldsymbol{x})$ is solved by QAs with short annealing time, the distribution of samples is more widespread around the ground state. If we consider that this distribution is the Boltzmann distribution and n samples $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ are generated by QAs, the thermal average is calculated as

$$\langle \mathcal{O}(\boldsymbol{x})\rangle_\beta^{\mathrm{QA}} = \frac{1}{n}\sum_{i=1}^{n} \mathcal{O}(\boldsymbol{x}_i). \tag{3}$$

However, the distribution by QAs deviates from the Boltzmann distribution, and the correction is needed.

B. Kernelized Stein discrepancy

Stein discrepancy can measure the difference between two distributions $p(\boldsymbol{x})$ and $q(\boldsymbol{x})$. Kernelized Stein discrepancy [20] is defined as

$$S(p, q) = \mathbb{E}_{\boldsymbol{x},\boldsymbol{x}'\sim q}[k_p(\boldsymbol{x}, \boldsymbol{x}')] = \mathrm{Tr}k_p(\boldsymbol{x}, \boldsymbol{x}')q(\boldsymbol{x})q(\boldsymbol{x}'), \tag{4}$$

where $k_p(\boldsymbol{x}, \boldsymbol{x}')$ is called the Stein kernel that depends on $p(\boldsymbol{x})$ and the base kernel

$$k(\boldsymbol{x}, \boldsymbol{x}') = \exp\left(-\frac{\sum_{i=1}^{d}\mathbb{I}\{x_i \neq x_i'\}}{d}\right). \tag{5}$$

The function of $\mathbb{I}\{x_i \neq x_i'\}$ shows 1 for $x_i \neq x_i'$ and 0 for the others, respectively. Let us define the difference operator as

$$\nabla_{\boldsymbol{x}} f(\boldsymbol{x}) = [f(\boldsymbol{x}) - f(\neg_1\boldsymbol{x}), \ldots, f(\boldsymbol{x}) - f(\neg_d\boldsymbol{x})], \tag{6}$$

where $\neg_i$ denotes the sign flip of the ith variable. In addition, the score function $\boldsymbol{s}_p(\boldsymbol{x}) \in \mathbb{R}^d$ is defined as

$$[\boldsymbol{s}_p(\boldsymbol{x})]_i = 1 - p(\neg_i\boldsymbol{x})/p(\boldsymbol{x}) \ \ (i = 1, \ldots, d).$$

The Stein kernel is then defined as

$$\begin{aligned} k_p(\boldsymbol{x}, \boldsymbol{x}') = {} & \boldsymbol{s}_p(\boldsymbol{x})^\top k(\boldsymbol{x}, \boldsymbol{x}')\boldsymbol{s}_p(\boldsymbol{x}') - \boldsymbol{s}_p(\boldsymbol{x})^\top \nabla_{\boldsymbol{x}'} k(\boldsymbol{x}, \boldsymbol{x}') \\ & - \boldsymbol{s}_p(\boldsymbol{x}')^\top \nabla_{\boldsymbol{x}} k(\boldsymbol{x}, \boldsymbol{x}') + \mathrm{Tr}[\nabla_{\boldsymbol{x},\boldsymbol{x}'} k(\boldsymbol{x}, \boldsymbol{x}')]. \end{aligned} \tag{7}$$

Notably, $k_p(\boldsymbol{x}, \boldsymbol{x}')$ depends on $p(\boldsymbol{x})$ only through the score functions. Here, $S(p, q)$ is always non-negative because $k_p(\boldsymbol{x}, \boldsymbol{x}')$ is a provably positive definite kernel if $k(\boldsymbol{x}, \boldsymbol{x}')$ is a positive definite kernel. In addition, $S(p, q) = 0$ holds if and only if $p = q$ [18].

To address the relationship between the probability distributions and their score functions, we consider the three-bit Ising model with random interactions. Figure 1 shows the Boltzmann distributions of the same Ising model with different temperatures as $\beta = 0.5$, 0.4, and 0.05 and their score functions. When the distributions are similar, their score functions are close, while when the distributions are dissimilar, their score functions diverge significantly. The Stein discrepancies are also shown in Fig. 1, and we have confirmed that the relationship between distributions can be correctly evaluated by the Stein discrepancy. Thus, we can understand that minimizing the difference in the score function minimizes the difference between distributions. Stein correction follows this idea of minimizing the difference in score functions.

C. Stein correction

In this study, we use the kernelized Stein discrepancy to correct the contributions of each sample obtained by QAs. Let $p(\boldsymbol{x})$ be the Boltzmann distribution $p(\boldsymbol{x}) \sim \exp[-\beta H_{\mathrm{Ising}}(\boldsymbol{x})]$ and $q(\boldsymbol{x})$ be an unknown distribution by QAs. Suppose that n samples are obtained from QAs. However, since $q(\boldsymbol{x})$ is unknown, the weight of each sample cannot be obtained. Thus, we assume that the weight of each sample is w_i instead of $q(\boldsymbol{x}_i)$, and the Stein discrepancy defined by Eq. (4) is approximated as

$$S(p, q) \simeq \sum_{i,j=1}^{n} w_i w_j k_p(\boldsymbol{x}_i, \boldsymbol{x}_j) = \boldsymbol{w}^\top K_p \boldsymbol{w}, \tag{8}$$

where K_p is the $n \times n$ matrix where elements are $k_p(\boldsymbol{x}_i, \boldsymbol{x}_j)$, which is called the Stein kernel matrix. In addition, $\boldsymbol{w}$ is the weight vector for n samples. In Stein correction [18], the weights are adjusted such that the Stein discrepancy is minimized,

$$\hat{\boldsymbol{w}} = \mathrm{argmin}_{\boldsymbol{w}}\left\{\boldsymbol{w}^\top K_p \boldsymbol{w} \text{ s.t. } w_i \geqslant 0, \sum_{i=1}^{n} w_i = 1\right\}. \tag{9}$$

Then, the samples weighted by $\hat{\boldsymbol{w}}$ approximate a Boltzmann distribution $p(\boldsymbol{x})$. Using the obtained weights, the thermal average of observables $\mathcal{O}(\boldsymbol{x})$ is approximately obtained by

$$\langle \mathcal{O}(\boldsymbol{x})\rangle_\beta^{\mathrm{SC}} = \sum_{i=1}^{n} \hat{w}_i \mathcal{O}(\boldsymbol{x}_i). \tag{10}$$

D. Fast Stein correction

A naive implementation of Stein correction requires $O(n^2)$ space and $O(n^3)$ time. We reduce the complexity by introducing a random feature map [21] and an exponentiated gradient descent [22]. Using the random feature map, the base kernel defined by Eq. (5) is approximated as the inner product $k(\boldsymbol{x}, \boldsymbol{x}') \approx \boldsymbol{\phi}(\boldsymbol{x})^\top \boldsymbol{\phi}(\boldsymbol{x}')$, where the feature map $\boldsymbol{\phi}(\boldsymbol{x}) : \{-1, 1\}^d \to \mathbb{R}^\ell$ is computed as follows. Let us draw ℓ samples $\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_\ell$ from

$$h(\boldsymbol{\omega}) = \prod_{i=1}^{d} \frac{1}{\pi\left(1 + \omega_i^2\right)}, \tag{11}$$

where $\boldsymbol{\omega}$ is the d-dimensional vector with each component ω_i. Also, $b_1, \ldots, b_\ell$ are sampled from the uniform distribution over $[0, 2\pi]$. Then, the feature map is defined as

$$\boldsymbol{\phi}(\boldsymbol{x}) = \left[z_{\boldsymbol{\omega}_1,b_1}\left(\frac{\boldsymbol{x}+\boldsymbol{1}}{2d}\right), \ldots, z_{\boldsymbol{\omega}_\ell,b_\ell}\left(\frac{\boldsymbol{x}+\boldsymbol{1}}{2d}\right)\right]^\top, \tag{12}$$

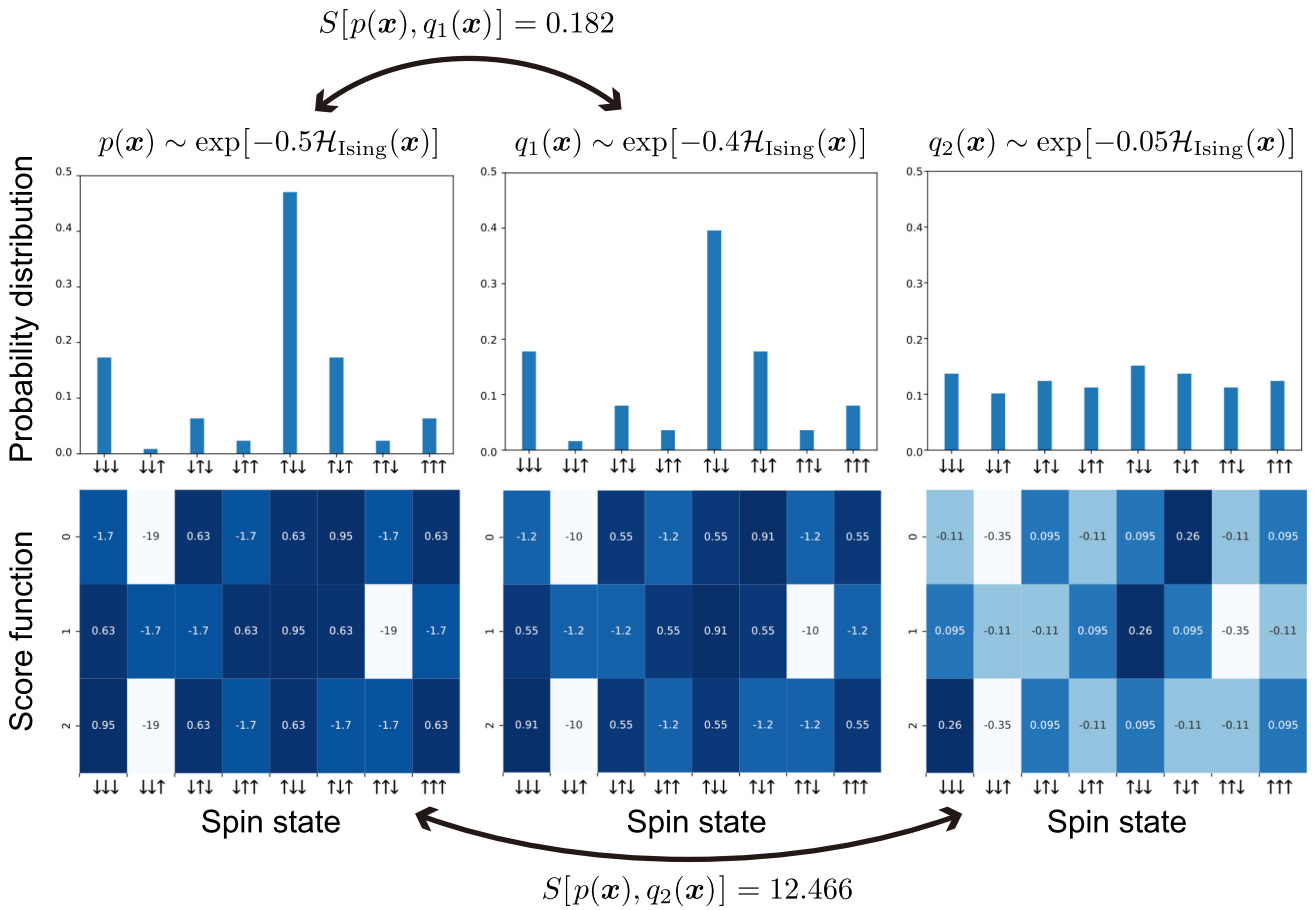

FIG. 1. Relationship between Boltzmann distributions and their score functions in three-bit Ising model with different temperatures as $\beta = 0.5$, 0.4, and 0.05. The Stein discrepancies between distributions are also shown.

where $z_{\omega,b}(\boldsymbol{x}) = \sqrt{\frac{2}{\ell}}\cos(\boldsymbol{\omega}^\top \boldsymbol{x} + b)$ and $\mathbf{1}$ is the d-dimensional vector of $(1, \ldots, 1)$. Since the Stein kernel is a linear function of the base kernel, it can also be approximated as $k_p(\boldsymbol{x}, \boldsymbol{x}') \approx \phi_p(\boldsymbol{x})^\top \phi_p(\boldsymbol{x})$, where $\phi_p(\boldsymbol{x})$ is the concatenation of the following vectors:

$$\theta_k(\boldsymbol{x}) = \frac{p(\neg_k \boldsymbol{x})}{p(\boldsymbol{x})}\phi(\boldsymbol{x}) - \phi(\neg_k \boldsymbol{x}) \quad (k = 1, \ldots, d). \tag{13}$$

Using the random feature map, the Stein discrepancy can be approximated as

$$S(p, q) \simeq \left\| \sum_{i=1}^{n} w_i \phi_p(\boldsymbol{x}_i) \right\|^2, \tag{14}$$

which is strictly positive. The optimization problem defined by Eq. (9) is a convex optimization problem with nonnegativity and normalization constraints. When the standard gradient descent algorithm is applied, the constraints are violated every time the parameters are updated. In this case, exponentiated gradient descent is known to work well [22], because constraint violation never happens. The update is described as

$$w_{t+1,i} = \frac{w_{t,i}\exp(-\eta[\nabla f(\boldsymbol{w}_t)]_i)}{Z_t}, \tag{15}$$

where

$$f(\boldsymbol{w}_t) = \sum_{i=1}^{n} w_{t,i}\phi_p(\boldsymbol{x}_i), \tag{16}$$

$$Z_t = \sum_{i=1}^{n} w_{t,i}\exp\{-\eta[\nabla f(\boldsymbol{w}_t)]_i\}, \tag{17}$$

and η is the learning rate. The modification shown above reduces the space requirement to $O(n\ell)$. Each update takes only $O(n)$ time, enabling us to deal with a large number of samples. The implementation of a fast Stein correction can be found on GitHub (see Ref. [23]). We have confirmed that approximately the same accuracy is obtained with the exact calculation of the Stein correction by our approximate algorithm (see the Supplemental Material, Fig. S1 [24]). Furthermore, we have measured the computation time as a function of the number of bits in the problem, d, with the number of samples fixed at 2 and the number of random feature maps $\ell = 5000$. The results are shown in Fig. S2 in Ref. [24], and the $O(d^2)$ computation time is required for the problem size.

III. RESULTS

In our experiments, we employ 16-bit Ising Hamiltonians proposed by Nelson *et al.* [10] for benchmarking. They are

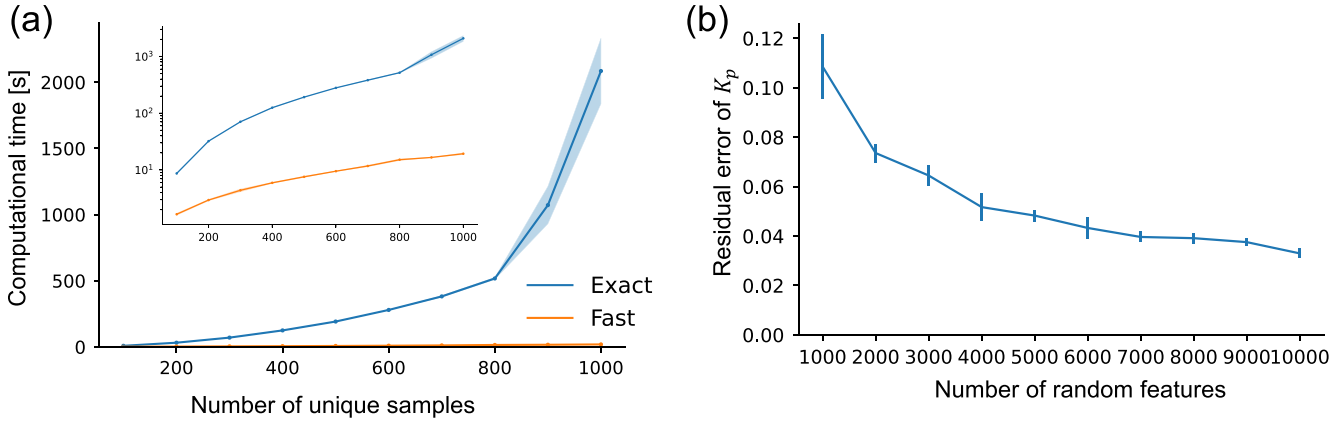

FIG. 2. Results for Stein correction on GSD_8. (a) Computational time of exact and fast Stein corrections depending on the number of samples n when $\ell = 5000$. The inset is the log scale figure. (b) Residual error of the Stein kernel matrix K_p against the number of random features ℓ when $n = 1000$. Five independent runs were performed, and the means and standard deviations are plotted as lines and error bars, respectively. The target Stein kernel matrix is common and random features are different towards five runs.

called GSD_X and the number X indicates the number of degenerated ground states. The Hamiltonians are designed to fit the chimera topology of D-Wave 2000Q systems, but here the samples are generated by Advantage System 6.2, because 2000Q is already out of service in their cloud platform. Hamiltonians are embedded in the Pegasus topology of Advantage using the MINORMINER PYTHON package [25].

First, we evaluate the efficiency and approximation error of fast Stein correction using GSD_8. Throughout the paper, the learning rate is $\eta = 10^{-5}$ and the number of updates is 3000 in the fast Stein correction. In addition, an exact correction was performed by CVXOPT [26]. Figure 2(a) shows the computational time of Stein correction depending on the number of samples when $\ell = 5000$. The fast correction was orders of magnitude faster than the exact one against the number of samples. Figure 2(b) shows the residual error of the Stein kernel matrix K_p, which is defined as $\|K_p - K_p^*\|/\|K_p^*\|$, depending on the number of features ℓ. Here, K_p^* is the exact matrix. The residual error by the random feature map decreases as the number of features is increased, showing that the more random features are preferred for accurate approximation. In the following analysis, calculations are performed with $\ell = 5000$, indicating a sufficiently better residual error and taking computational time into account. Note that residual errors depending on the system size d are shown in Fig. S3 in Ref. [24], and when the number of features is larger than 20 000, the residual error becomes efficiently small up to 1024 bits.

Next, we observe how fast Stein correction improves the accuracy of estimating observables on GSD_8, GSD_38, and GSD_F_6. For GSD_F_6, the finite fields h_i are imposed. The calculated thermal averages of observables are the internal energy $E(\beta)$, the magnetic susceptibility $\chi(\beta)$, and the Binder cumulant $U_4(\beta)$, defined by

$$E(\beta) = \langle H_{\mathrm{Ising}}(\boldsymbol{x})\rangle_\beta, \tag{18}$$

$$\chi(\beta) = \beta\left\langle\left(\sum_{i=1}^d x_i\right)^2\right\rangle_\beta, \tag{19}$$

$$U_4(\beta) = 1 - \left\langle\left(\sum_{i=1}^d x_i\right)^4\right\rangle_\beta / 3\left\langle\left(\sum_{i=1}^d x_i\right)^2\right\rangle_\beta^2, \tag{20}$$

where d is the system size. For each observable, we compute the residual error as $\|y - y^*\|/y^*$, where y is the thermal average of observables calculated by Eqs. (3) and (10) and y^* is the exact value computed via brute-force enumeration. For each β, 10 000 samples are generated by a D-Wave QA by setting num_reads = 10 000 with an annealing time of 5 μs. For this setting, when β is increased, the number of unique samples is decreased in 10 000 (see Fig. S4 in Ref. [24]). In addition, the Metropolis method with single-spin flip, a basic MCMC method, has also been applied. In our implementation, the initial state is randomly generated at each temperature, and the MCMC calculation is quenched at the target temperature. In the MCMC calculation, we restart the Markov chain when the length of the chain reaches 2000. This corresponds to the fact that a single MCMC run consists of 5 independent Markov chains. To match the number of QA samples, the first 1500 samples are used for burn-in and the remaining 500 samples are used for estimation. Computing time comparisons for each inverse temperature between Stein-corrected QAs and MCMC calculation are shown in Fig. S5 in Ref. [24].

The results over a range of inverse temperatures are shown in Fig. 3. For each case, five independent runs are conducted and the mean values of residual error are evaluated. The accuracy of fast Stein correction outperforms the original QA samples in all cases, showing the effectiveness of our approach. The temperature dependence for internal energy behaves nonmonotonically, and there is a temperature at which the residual error reaches a minimum. If appropriate samples are obtained from D-Wave QAs for each temperature, the results by Stein correction will be more accurate. Thus, we consider that the temperature with the minimum error is close to the effective temperature of the D-Wave QA, which is considered to give the best sampling. On the other hand, the accuracies by the Stein correction and QA samples are low at lower inverse temperatures (high temperatures) in the internal energy. Even if high temperatures are set, the low-energy states are obtained by QAs and these are not important samples in the thermal averages. Thus, if the temperature is far from the effective temperature of D-Wave QAs, the accuracy can be expected to improve by adding appropriate non-QA

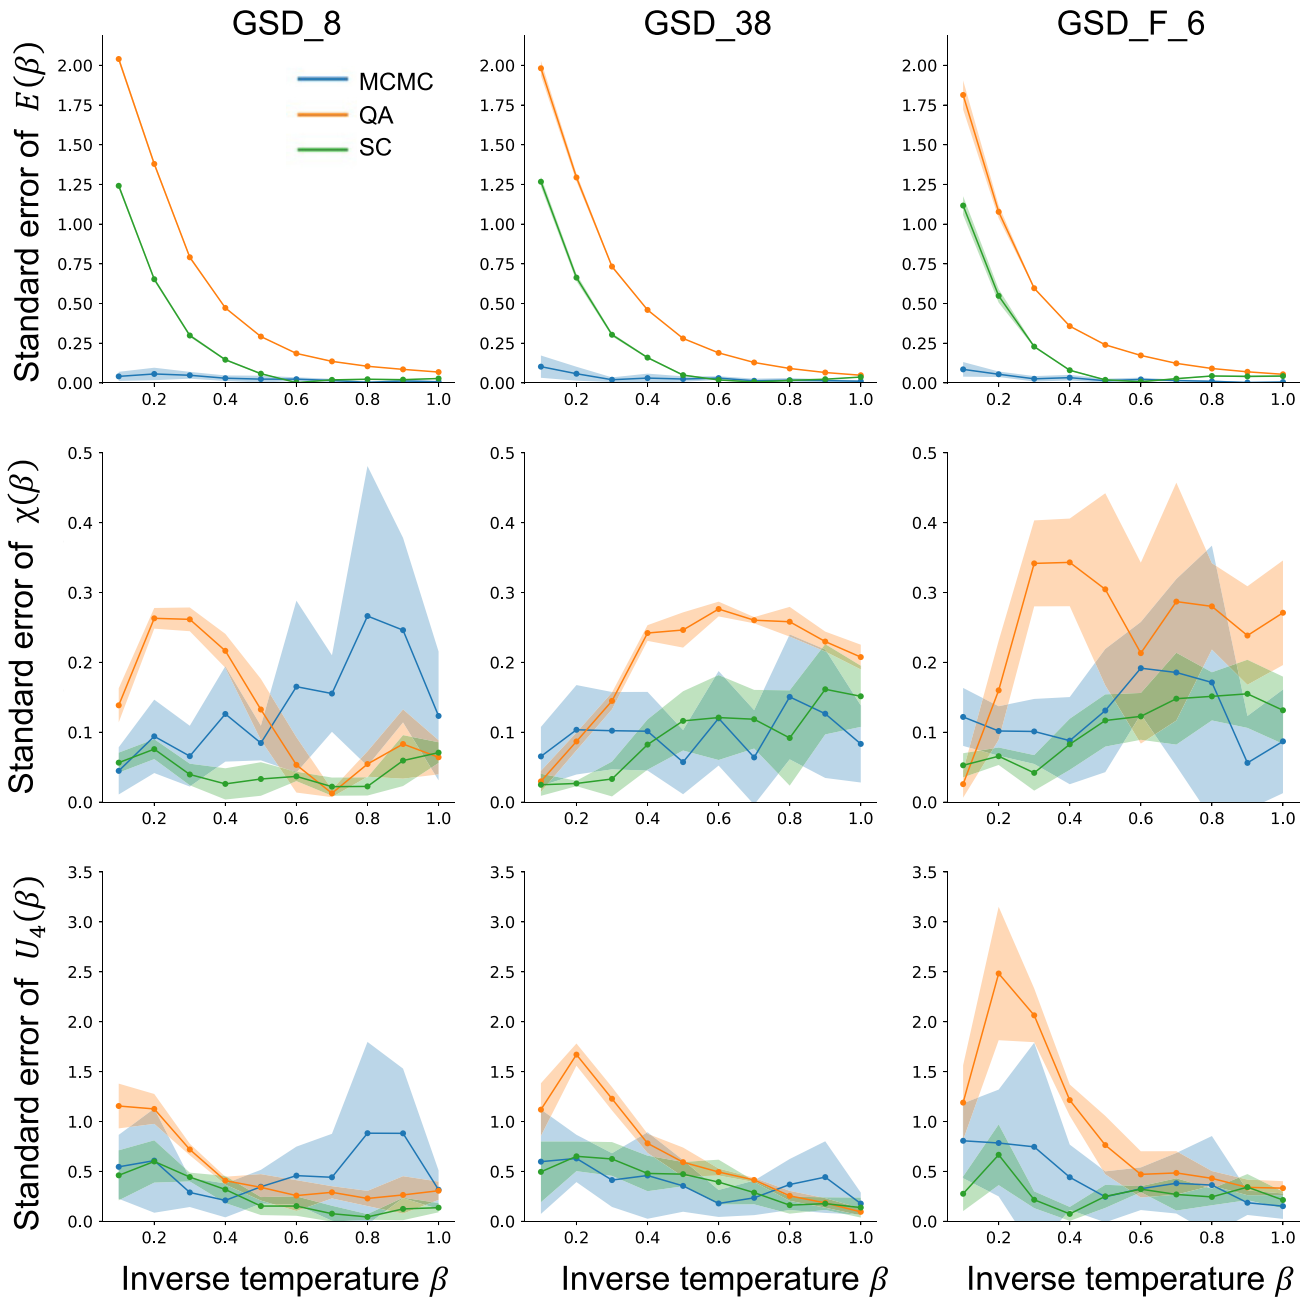

FIG. 3. Residual errors of internal energy $E(\beta)$, magnetic susceptibility $\chi(\beta)$, and Binder cumulant $U_4(\beta)$. The thermal averages were calculated by MCMC, a naive average by quantum annealers (QA) and Stein correction (SC), respectively. The number of samples was fixed as $n = 10\,000$, and the results depending on n are shown in Fig. S6 in Ref. [24] when $\beta = 0.4$. Five independent runs were performed, and the means and standard deviations are plotted as lines and shaded areas, respectively.

samples. In addition, the residual errors depending on the number of samples when $\beta = 0.4$ are summarized in Fig. S6 in Ref. [24], resulting in the case that the error does not always decrease with more samples. One reason is that the number of unique samples is smaller and smaller as the temperature is decreased. In fact, many samples are the same in the QA samples (see Fig. S4 in Ref. [24]). In particular, errors on magnetization and the Binder cumulant are not decreased as the number of samples is increased. This suggests that expectation calculation of magnetization and the Binder cumulant is essentially more difficult and sensitive to samples compared to the internal energy.

Furthermore, the error by fast Stein correction was comparable to that by MCMC except for the internal energy at high-temperature regions. These results show that fast Stein correction has the potential to expand the applicability range of quantum annealers significantly and may replace MCMC in diverse tasks of discrete sampling.

Since the 16-bit system is used, the exact distribution depending on the temperature can be obtained by brute-force

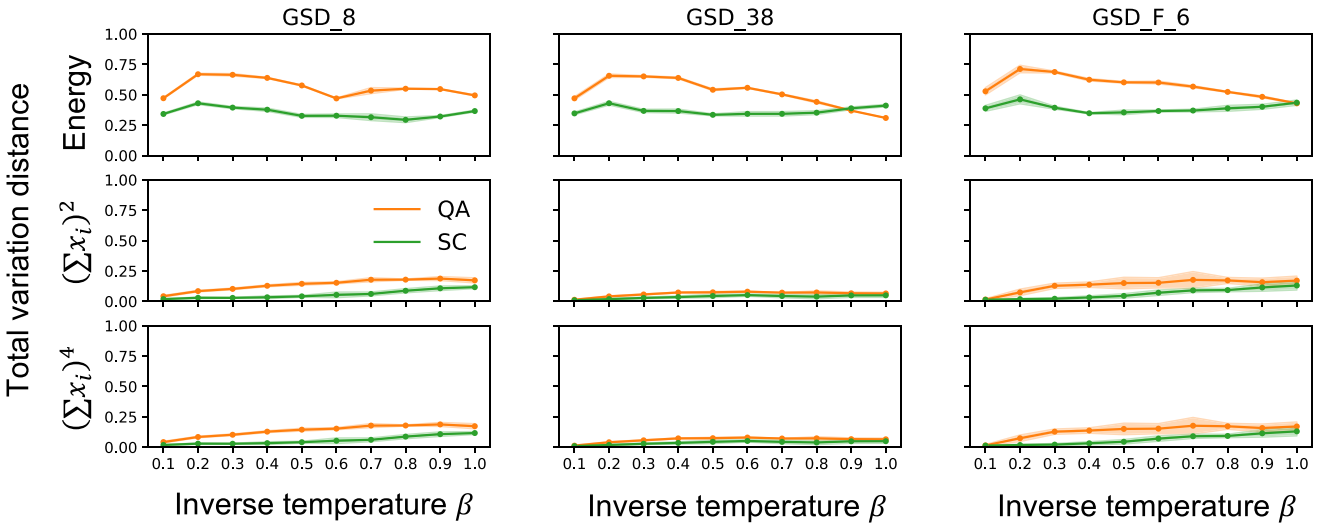

FIG. 4. Total variation distance between exact distributions and Stein-corrected distributions for energies of states $(\sum_{i=1}^{d} x_i)^2$ and $(\sum_{i=1}^{d} x_i)^4$, respectively. The results by QA samples are also shown.

technique. Thus, we can compare the exact distributions and the Stein-corrected distribution directly, and the total variation distance depending on the temperature is summarized in Fig. 4. The energies of states $(\sum_{i=1}^{d} x_i)^2$ and $(\sum_{i=1}^{d} x_i)^4$ are the target observables. By comparing with the results of QAs, we confirm that the Stein correction can fit the distributions to multiple observables simultaneously.

IV. CONCLUSION

We demonstrated that fast Stein correction is a helpful companion of quantum annealers and fundamentally enhances their usability. The proposed method will play a role as an efficient postprocessing for quantum annealing [11] to achieve temperature-dependent Boltzmann sampling. Since our approach is a general method that can obtain the desired distribution given arbitrary samples, it can be applied to other software- and hardware-based sampling methods. The advantage of QA samples is that they are not locally concentrated, whereas MCMC samples have difficulty covering the whole space. We believe that our proposed method combining QAs and Stein corrections is useful in many kinds of applications where MCMC does not work efficiently and high-quality samples cannot be obtained by random sampling. For example, our method is effective for constrained optimization problems. When dealing with problems with equality constraints, random walk does not work efficiently because some states satisfying constraints are not reachable from each other by a single spin flip. Generating random samples by MCMC is thus difficult in constrained problems. On the other hand, QAs can consider equality constraints by using appropriate procedures when designing the Ising model [27]. Therefore, high-quality samples can be obtained even when such constraints are included.

Although Stein correction cannot bring the distributional error to zero, it would be particularly useful to sample from highly constrained spaces [28], where global mixing by MCMC is extremely hard. Our future work involves the application of our method to machine learning and statistical physics and other highly scalable Ising machines such as coherent Ising machines [29] and GPU-based algorithms [30]. It can also be applied to general quantum computers including NISQ, where distributional error is unavoidable due to environmental noise [31].

ACKNOWLEDGMENTS

R.S. thanks participants at AQC2023 in Albuquerque, New Mexico, for fruitful discussions. This work is supported by AIP Kasoku under Grant No. JPMJCR21U2, JST CREST under Grant No. JPMJCR21O2, JST ERATO under Grant No. JPMJER1903, KAKENHI under Grant No. 19H05819, and MEXT under Grant No. JPMXP1122712807.

[1] K. Binder and A. P. Young, Spin glasses: Experimental facts, theoretical concepts, and open questions, Rev. Mod. Phys. **58**, 801 (1986).

[2] K. Binder and E. Luijten, Monte Carlo tests of renormalization-group predictions for critical phenomena in Ising models, Phys. Rep. **344**, 179 (2001).

[3] N. Zhang, S. Ding, J. Zhang, and Y. Xue, An overview on restricted Boltzmann machines, Neurocomputing **275**, 1186 (2018).

[4] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, Restricted Boltzmann machines in quantum physics, Nat. Phys. **15**, 887 (2019).

[5] A. Doucet, N. De Freitas, N. J. Gordon *et al.*, *Sequential Monte Carlo methods in practice* (Springer, New York, 2001).

[6] K. Hukushima and K. Nemoto, Exchange Monte Carlo method and application to spin glass simulations, J. Phys. Soc. Jpn. **65**, 1604 (1996).

[7] Y. Iba, Extended ensemble Monte Carlo, Int. J. Mod. Phys. C **12**, 623 (2001).

[8] K. Hukushima and Y. Iba, Population annealing and its application to a spin glass, AIP Conf. Proc. **690**, 200 (2003).

[9] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom *et al.*, Quantum annealing with manufactured spins, Nature (London) **473**, 194 (2011).

[10] J. Nelson, M. Vuffray, A. Y. Lokhov, T. Albash, and C. Coffrin, High-Quality thermal Gibbs sampling with quantum annealing hardware, Phys. Rev. Appl. **17**, 044046 (2022).

[11] J. Raymond, S. Yarkoni, and E. Andriyash, Global warming: Temperature estimation in annealers, Front. ICT **3**, 23 (2016).

[12] J. Marshall, D. Venturelli, I. Hen, and E. G. Rieffel, Power of pausing: Advancing understanding of thermalization in experimental quantum annealers, Phys. Rev. Appl. **11**, 044083 (2019).

[13] J. Marshall, E. G. Rieffel, and I. Hen, Thermalization, freeze-out, and noise: Deciphering experimental quantum annealers, Phys. Rev. Appl. **8**, 064025 (2017).

[14] R. Sandt and R. Spatschek, Efficient low temperature Monte Carlo sampling using quantum annealing, Sci. Rep. **13**, 6754 (2023).

[15] M. Vuffray, C. Coffrin, Y. A. Kharkov, and A. Y. Lokhov, Programmable quantum annealers as noisy Gibbs samplers, PRX Quantum **3**, 020317 (2022).

[16] Y. Matsuda, H. Nishimori, and H. G. Katzgraber, Ground-state statistics from annealing algorithms: quantum versus classical approaches, New J. Phys. **11**, 073021 (2009).

[17] S. T. Tokdar and R. E. Kass, Importance sampling: A review, Wiley Interdiscip. Rev. Comput. Stat. **2**, 54 (2010).

[18] Q. Liu and J. Lee, Black-box Importance Sampling, in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research (PMLR, 2017), Vol. 54, pp. 952–961.

[19] L. Hodgkinson, R. Salomone, and F. Roosta, The reproducing Stein kernel approach for post-hoc corrected sampling, arXiv:2001.09266.

[20] J. Yang, Q. Liu, V. Rao, and J. Neville, Goodness-of-fit testing for discrete distributions via Stein discrepancy, in *Proceedings of the 35th International Conference on Machine Learning*, Proceedings of Machine Learning Research (PMLR, 2018), Vol. 80, pp. 5561–5570.

[21] A. Rahimi and B. Recht, Random features for large-scale kernel machines, *Proceedings of the 20th International Conference on Neural Information Processing Systems* (Curran Associates, Inc., 2007), p. 1177.

[22] J. Kivinen and M. K. Warmuth, Exponentiated gradient versus gradient descent for linear predictors, Inf. Comput. **132**, 1 (1997).

[23] GitHub, https://github.com/tsudalab/fast-stein-correction.

[24] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevResearch.6.043050 for Fig. S1: Residual errors by exact and fast calculations of Stein correction; Fig. S2: Computational time depending on the problem size; Fig. S3: Residual error depending on the number of random features; Fig. S4: Number of unique samples depending on the target temperature; Fig. S5: Elapsed time between Stein-correction and Metropolis algorithm; Fig. S6: Residual errors of physical quantities depending on the number of samples.

[25] Dwavesystems/minorminer, https://github.com/dwavesystems/minorminer, accessed: 2023-08-03.

[26] M. S. Andersen, J. Dahl, L. Vandenberghe *et al.*, CVXOPT: A Python package for convex optimization, https://cvxopt.org/ 54 (2013).

[27] I. Hen, Equation planting: A tool for benchmarking Ising machines, Phys. Rev. Appl. **12**, 011003(R) (2019).

[28] S. Aoki, H. Hara, and A. Takemura, *Markov Bases in Algebraic Statistics* (Springer, Berlin, 2012).

[29] T. Inagaki, Y. Haribara, K. Igarashi, T. Sonobe, S. Tamate, T. Honjo, A. Marandi, P. L. McMahon, T. Umeki, K. Enbutsu, O. Tadanaga, H. Takenouchi, K. Aihara, K.-I. Kawarabayashi, K. Inoue, S. Utsunomiya, and H. Takesue, A coherent Ising machine for 2000-node optimization problems, Science **354**, 603 (2016).

[30] Z. Mao, Y. Matsuda, R. Tamura, and K. Tsuda, Chemical design with GPU-based Ising machines, Digital Discovery **2**, 1098 (2023).

[31] E. Pelofske, J. Golden, A. Bärtschi, D. O'Malley, and S. Eidenbenz, Sampling on NISQ devices: "Who's the fairest one of all?", in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, New York, 2021), pp. 207–217.