

Performance-guaranteed regularization in maximum likelihood method: Gauge symmetry in Kullback-Leibler divergence

Akihisa Ichiki 

Department of Applied Mathematics, Faculty of Science, Fukuoka University, 8-19-1, Nanakuma, Jonan-ku,
Fukuoka City 814-0180, Japan



(Received 29 March 2023; accepted 29 September 2023; published 19 October 2023)

The maximum likelihood method is the best-known method for estimating the probabilities behind the data. However, the conventional method obtains the probability model closest to the empirical distribution, resulting in overfitting. Then regularization methods prevent the model from being excessively close to the wrong probability, but little is known systematically about their performance. The idea of regularization is similar to error-correcting codes, which obtain optimal decoding by mixing suboptimal solutions with an incorrectly received code. The optimal decoding in error-correcting codes is achieved based on gauge symmetry. We propose a theoretically guaranteed regularization in the maximum likelihood method by focusing on a gauge symmetry in Kullback-Leibler divergence. In our approach, we obtain the optimal model without the need to search for hyperparameters frequently appearing in regularization.

DOI: [10.1103/PhysRevE.108.044134](https://doi.org/10.1103/PhysRevE.108.044134)

I. INTRODUCTION

The maximum likelihood method is often employed to estimate the ground truth distribution $P_{\text{GT}}(x)$ ($x \in \mathbb{R}^D$) from the given data that generate them. The basic criterion is to find a probability model $P(x)$ that minimizes the Kullback-Leibler divergence (KL divergence) between the probability model P and the underlying distribution P_{GT} [1,2]:

$$D[P_{\text{GT}}, P] := \int dx P_{\text{GT}}(x) \ln \frac{P_{\text{GT}}(x)}{P(x)}. \quad (1)$$

Since the ground truth P_{GT} is unknown and to be estimated, the well-known maximum likelihood method employs the empirical distribution P_{emp}^{ξ} as an approximation of P_{GT} . However, the probability model P that minimizes the KL divergence $D[P_{\text{emp}}^{\xi}, P]$ results in the model far from P_{GT} because of the difference between the empirical distribution P_{emp}^{ξ} and the ground truth P_{GT} . This phenomenon is well-known in machine learning as overfitting [3,4]. To avoid overfitting, regularization is applied to prevent $D[P_{\text{emp}}^{\xi}, P]$ from becoming too small [5–8]. Note that much of the success of regularization is backed by empirical evidence, and little is guaranteed theoretically on its performance. For the maximum likelihood estimation as model selection, Akaike's information criterion (AIC) gives a theoretically guaranteed regularization [9–11]. The method proposed in the present paper shares the same root with AIC. While AIC is defined based on an unbiased estimator of the KL divergence between the model and the ground truth, our method does not require such a quantity explicitly.

Regularization to prevent overfitting is similar to error-correcting codes [12–14]. In error-correcting codes, a sender

encodes the message redundantly, and a receiver decodes it by finding the solution to an optimization problem corresponding to the received code. However, since the communication channel is noisy, the received code contains errors. Thus, the optimization problem corresponding to the received code does not give correct decoding. The receiver utilizes a mixture of nonoptimal solutions to the optimization problem corresponding to the received code to decode the correct message [15,16]. In correspondence with physics, the optimization problem corresponds to the search for the ground state for the Hamiltonian, which the received code corresponds to. Since the Hamiltonian corresponding to the received code is incorrect, the correct message corresponds to the excited state for the incorrect Hamiltonian. In this situation, a mixture of excited states at finite temperature is more appropriate than only the ground state at zero temperature [17–20]. The condition for maximum restoration in error-correcting codes corresponds to the Nishimori condition known in the field of spin glasses [21–23]. In the case of maximum likelihood estimation, the ground truth P_{GT} influences the empirical distribution, but the two are different. Thus, it may be better to mix a probability model with a high value of the KL divergence to the ground truth than to select an optimal model. In this idea, the KL divergence plays the role of the Hamiltonian in maximum likelihood estimation.

The structure of the present paper is as follows: in the next section, we introduce a gauge transformation that keeps the KL divergence to obtain the Nishimori condition for probability model estimation. In Sec. III, we show that the probability model given by the Nishimori condition minimizes the KL divergence with the ground truth P_{GT} in the sense of expectation. The conventional maximum likelihood estimation selects an optimal probability model with the highest explanatory power for the given data from the prepared set of models. Our method requires a change in the concept of model

*ichiki@fukuoka-u.ac.jp

selection from that in the conventional maximum likelihood method. We discuss the two possible changes in Sec. IV. The first is to select a model probabilistically, and the second is to construct a new model consisting of prepared probability models. Section V presents two simple examples. The first is a model selection from a set formed by only two probability models. The second example is the set of probability models of the entire one-dimensional normal distributions. With prior knowledge that the ground truth P_{GT} is a normal distribution, it is intuitively expected that the best probability model obeys normal distribution. However, according to the latter interpretation described in Sec. IV, we show that the best probability model given by our method is not a normal distribution. Section VI discusses improvements to the set of models. Maximum likelihood estimation requires a set of prepared probability models. Since the performance of maximum likelihood estimation depends on this set, improvement of the set is preferred. Section VII discusses the relation between our proposed method and Bayesian updating, which is standard for estimating probability distribution based on given data. We discuss the sequential extension of the set. The last section is devoted to a conclusion and future perspectives.

II. GAUGE INVARIANCE IN KL DIVERGENCE

In our framework, the KL divergence between the empirical distribution and a probability model corresponds to a Hamiltonian in the theory of spin glasses. Under the Nishimori condition [21,23] for spin glasses, the exact internal energy is calculated. To derive the Nishimori condition, which plays a central role in error-correcting codes and also in our discussion, let us consider the expected value of the KL divergence between the empirical distribution and a probability model corresponding to the internal energy in spin glasses:

$$\langle \langle D[P_{\text{emp}}^{\xi}, P] \rangle \rangle_{\beta, \xi} \Big|_{\text{data}}, \quad (2)$$

where $\langle \cdot \rangle_{\text{data}}$ denotes the expected value with respect to the sample data $\xi = \{\xi_i\}_{i=1, \dots, N}$, i.e.,

$$\langle \cdot \rangle_{\text{data}} := \int \cdot \prod_{i=1}^N [d\xi_i P_{\text{GT}}(\xi_i)]. \quad (3)$$

The data ξ are assumed to be sampled independently and identically from the ground truth P_{GT} . In addition, $\langle \cdot \rangle_{\beta, \xi}$ represents the Gibbs–Boltzmann-type weighted average with inverse temperature β . The Hamiltonian in our framework is the KL divergence between the empirical distribution $P_{\text{emp}}^{\xi}(x) := \frac{1}{N} \sum_{i=1}^N \delta(x - \xi_i)$ and a probability model P :

$$\langle \cdot \rangle_{\beta, \xi} := \frac{\sum_{P \in \mathcal{M}} \cdot \exp\{-\beta D[P_{\text{emp}}^{\xi}, P]\}}{\sum_{P \in \mathcal{M}} \exp\{-\beta D[P_{\text{emp}}^{\xi}, P]\}}, \quad (4)$$

where $\sum_{P \in \mathcal{M}} \cdot$ is the sum over all models P in the set \mathcal{M} of probability models under consideration.

In the argument of the Nishimori condition in spin glass theory, the gauge invariance plays an important role. The Hamiltonian is invariant under the simultaneous transformation on the thermally fluctuating spins as stochastic variables and on the coupling constants as quenched random variables. In our case of the mixture of various probability models, the

model P corresponds to a stochastic variable. In contrast, since the data ξ are generated randomly from P_{GT} and fixed once given, they can be considered quenched random. Let us introduce the following gauge transformation:

$$P(x) \rightarrow P_f(x) := \det[\nabla^T f(x)] P(f(x)), \quad (5)$$

$$\xi_i \rightarrow \xi_i^f := f^{-1}(\xi_i), \quad (6)$$

where f is an arbitrary bijective map $\mathbb{R}^D \rightarrow \mathbb{R}^D$ and f^{-1} is its inverse map. It is straightforward to show that $P_f(x)$ given by Eq. (5) is indeed a density function. Since the empirical distribution is transformed under the gauge transformation as

$$\begin{aligned} P_{\text{emp}}^{\xi^f}(x) &= \frac{1}{N} \sum_{i=1}^N \delta[x - f^{-1}(\xi_i)] \\ &= \det[\nabla^T f(x)] P_{\text{emp}}^{\xi}(f(x)), \end{aligned} \quad (7)$$

we find

$$\frac{P_{\text{emp}}^{\xi^f}(x)}{P_f(x)} = \frac{P_{\text{emp}}^{\xi}(f(x))}{P(f(x))}. \quad (8)$$

In Eq. (7), we have used $\delta\{f[x - f^{-1}(\xi_i)]\} = 0$ when $f(x) \neq \xi_i$ since f is a bijective map. In addition, $\nabla^T f(x)$ represents the Jacobian of $f(x)$. By the variable transformation $y = f(x)$, the KL divergence is shown to be invariant under the gauge transformation:

$$\begin{aligned} D[P_{\text{emp}}^{\xi^f}, P_f] &= \int dx \det[\nabla^T f(x)] P_{\text{emp}}^{\xi}(f(x)) \ln \frac{P_{\text{emp}}^{\xi}(f(x))}{P(f(x))} \\ &= \int dy P_{\text{emp}}^{\xi}(y) \ln \frac{P_{\text{emp}}^{\xi}(y)}{P(y)} = D[P_{\text{emp}}^{\xi}, P]. \end{aligned} \quad (9)$$

On the other hand, the probability of the ground truth for infinitesimal volume element $d\xi$ varies under the gauge transformation as

$$\begin{aligned} P_{\text{GT}}(\xi) d\xi &\rightarrow P_{\text{GT}}(\xi^f) d\xi^f \\ &= \det[\nabla^T f^{-1}(\xi)] P_{\text{GT}}(f^{-1}(\xi)) d\xi \\ &:= P_{\text{GT}}^f(\xi) d\xi. \end{aligned} \quad (10)$$

Note that the following theorem holds [24–26]: for any pair $[p(x), q(x)]$ of density functions in D dimensions, there exists a convex function ψ satisfying

$$p(x) = [\det \nabla \nabla^T \psi(x)] q(\nabla \psi(x)), \quad (11)$$

where $\nabla \nabla^T \psi(x)$ denotes the Hessian of $\psi(x)$. The function ψ is unique up to an arbitrary additive constant. Thus, the target density $q(y)$ can be obtained from any density $p(x)$ by applying an appropriate variable transformation $y = \nabla \psi(x)$, and such a variable transformation is unique. In our framework, the gauge transformation with $f^{-1}(x) = \nabla \psi(x)$ uniquely determines P_{GT}^f connected to the ground truth P_{GT} . Using the invariance of the KL divergence under the gauge

transformation, we find

$$\begin{aligned} \langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{\beta, \xi | \text{data}} &= \int \prod_{i=1}^N [d\xi_i P_{\text{GT}}(\xi_i)] \frac{D_{\mathcal{M}, I}^{\xi, \beta}}{Z_{\mathcal{M}, I}^{\xi, \beta}} \\ &= \int \prod_{i=1}^N [d\xi_i^f P_{\text{GT}}(\xi_i^f)] \frac{D_{\mathcal{M}, f}^{\xi, \beta}}{Z_{\mathcal{M}, f}^{\xi, \beta}} \\ &= \int \prod_{i=1}^N [d\xi_i^f P_{\text{GT}}(\xi_i^f)] \frac{D_{\mathcal{M}, I}^{\xi, \beta}}{Z_{\mathcal{M}, I}^{\xi, \beta}}, \end{aligned} \quad (12)$$

where $D_{\mathcal{M}, f}^{\xi, \beta} := \sum_{P_f \in \mathcal{M}_f} D[P_{\text{emp}}^f, P_f] e^{-\beta D[P_{\text{emp}}^f, P_f]}$ and $Z_{\mathcal{M}, f}^{\xi, \beta} := \sum_{P_f \in \mathcal{M}_f} e^{-\beta D[P_{\text{emp}}^f, P_f]}$ with the notation \mathcal{M}_f denoting the image of \mathcal{M} by f . The relations $P_{\text{emp}}^f = P_{\text{emp}}^\xi$, $P_f = P$, and $\mathcal{M}_f = \mathcal{M}$ have been used for an identity map I . In the last equality in Eq. (12), the expression has been rewritten as a sum over $P \in \mathcal{M}$ since each $P_f \in \mathcal{M}_f$ corresponds to an appropriate $P \in \mathcal{M}$ by the gauge transformation. In addition, since an arbitrary density $P_{\text{GT}}^f(\xi)$ can be represented as

$$P_{\text{GT}}^f(\xi) = \exp \left[\int dx \delta(x - \xi) \ln P_{\text{GT}}^f(x) \right], \quad (13)$$

$$\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{\beta, \xi | \text{data}} = \frac{1}{N_{\mathcal{F}}} \int \prod_{i=1}^N d\xi_i \left[\sum_{f \in \mathcal{F}} E_{P_{\text{GT}}^f}^{\xi, N} \right] \frac{\sum_{P \in \mathcal{M}} D[P_{\text{emp}}^\xi, P] E_P^{\xi, \beta}}{\sum_{P \in \mathcal{M}} E_P^{\xi, \beta}}. \quad (17)$$

Using the correspondence between f and P_{GT}^f , the sum for f becomes the sum for P_{GT}^f . Then we have

$$\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{\beta, \xi | \text{data}} = \frac{1}{N_{\mathcal{M}_{\text{GT}}^{\mathcal{F}}}} \int \prod_{i=1}^N d\xi_i \left[\sum_{P \in \mathcal{M}} D[P_{\text{emp}}^\xi, P] E_P^{\xi, \beta} \right] \frac{\sum_{P \in \mathcal{M}_{\text{GT}}^{\mathcal{F}}} E_P^{\xi, N}}{\sum_{P \in \mathcal{M}} E_P^{\xi, \beta}}, \quad (18)$$

where $\mathcal{M}_{\text{GT}}^{\mathcal{F}}$ is the set composed of all of P_{GT}^f given by the bijective maps $f \in \mathcal{F}$, and $N_{\mathcal{M}_{\text{GT}}^{\mathcal{F}}}$ is the number of its elements (or volume in the case of a continuum). Note that $N_{\mathcal{F}} = N_{\mathcal{M}_{\text{GT}}^{\mathcal{F}}}$ holds when \mathcal{F} is a discrete set. The ratio of Boltzmann weights appearing on the right-hand side in Eq. (18) becomes unity when $\beta = N$ and $\mathcal{M}_{\text{GT}}^{\mathcal{F}} = \mathcal{M}$. The condition for the unity in the ratio is called the Nishimori condition hereafter.

The condition for $\mathcal{M}_{\text{GT}}^{\mathcal{F}} = \mathcal{M}$ is given as follows. We can use an arbitrary bijective map f in the gauge transformation. Since $P_{\text{GT}}^f(x) = \det[\nabla^T f^{-1}(x)] P_{\text{GT}}(f^{-1}(x))$ and, for any density function Q , there exists a unique bijective map g satisfying $Q(x) = \det[\nabla^T g(x)] P_{\text{GT}}(g(x))$, we can find a unique bijective map f satisfying $P_{\text{GT}}^f(x) = P(x)$ for each model $P \in \mathcal{M}$. Thus, assuming that \mathcal{F} is the set consisting of all bijective maps f satisfying $P_{\text{GT}}^f(x) = P(x) \in \mathcal{M}$, $\mathcal{M}_{\text{GT}}^{\mathcal{F}} = \mathcal{M}$ with $N_{\mathcal{F}} = N_{\mathcal{M}}$ holds, where $N_{\mathcal{M}}$ is the number of elements of \mathcal{M} (or volume in the case of a continuum). Since P_{GT} is unknown, to give a concrete set of the bijective maps \mathcal{F} is impossible. However, such a set exists, and a concrete expression of \mathcal{F} is not required. Therefore, for an arbitrary set of probability models \mathcal{M} , there exists an appropriate set of bijective maps \mathcal{F} that allows the ratio of Boltzmann weights appearing on the

and the formal expression

$$\prod_{i=1}^N e^{\int dy \delta(y - \xi_i) \ln P_{\text{GT}}^f(y)} = e^{N \int dy P_{\text{emp}}^\xi(y) \ln P_{\text{GT}}^f(y)} \quad (14)$$

holds, Eq. (12) can be rewritten through Eqs. (10), (13), and (14) as

$$\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{\beta, \xi | \text{data}} = \int \prod_{i=1}^N d\xi_i E_{P_{\text{GT}}^f}^{\xi, N} \frac{D_{\mathcal{M}, I}^{\xi, \beta}}{Z_{\mathcal{M}, I}^{\xi, \beta}}, \quad (15)$$

where $E_P^{\xi, \beta} := e^{\beta \int dy P_{\text{emp}}^\xi(y) \ln P(y)}$. Since $\int dx P_{\text{emp}}^\xi(x) \ln P_{\text{emp}}^\xi(x)$ in the definition of the KL divergence is independent of the model P , we have

$$\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{\beta, \xi | \text{data}} = \int \prod_{i=1}^N d\xi_i E_{P_{\text{GT}}^f}^{\xi, N} \frac{\sum_{P \in \mathcal{M}} D[P_{\text{emp}}^\xi, P] E_P^{\xi, \beta}}{\sum_{P \in \mathcal{M}} E_P^{\xi, \beta}}. \quad (16)$$

Since $\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{\beta, \xi | \text{data}}$ is invariant under the gauge transformation, the result of Eq. (16) is independent of f . Thus, the result is the same as the average over all elements f of an arbitrary set of bijective maps \mathcal{F} . Denoting the number (or volume in the case of a continuum) of the bijective maps by $N_{\mathcal{F}}$, we obtain

right-hand side of Eq. (18) to become unity. The Nishimori condition is always satisfied by $\beta = N$ and $\mathcal{M}_{\text{GT}}^{\mathcal{F}} = \mathcal{M}$. In this case, Eq. (18) yields the simple form as

$$\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{N, \xi | \text{data}} = \frac{1}{N_{\mathcal{M}}} \sum_{P \in \mathcal{M}} \int \prod_{i=1}^N [d\xi_i P(\xi_i)] D[P_{\text{emp}}^\xi, P]. \quad (19)$$

From this expression, $\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{N, \xi | \text{data}}$ is independent of P_{GT} and is given as the average over $P \in \mathcal{M}$ of the expected KL divergence between the model P and the empirical distribution given by the realizations of the data $\{\xi_i\}_{i=1, \dots, N}$ generated from P . In other words, the expectation for the KL divergence is calculated only from the probability models, in principle, on the Nishimori condition. It is consistent with the fact that the internal energy is calculated exactly on the Nishimori condition in the framework of spin glass theory [20].

In the framework of AIC, the unbiased estimator for the KL divergence between the empirical distribution P_{emp}^ξ and the probability model P is evaluated specifically [9, 11]. Correspondingly, it has been shown that $\langle\langle D[P_{\text{emp}}^\xi, P] \rangle\rangle_{N, \xi | \text{data}}$ can

be calculated, in principle, in our case. However, as seen in the next section, such a concrete expression is not required for the best estimation of P_{GT} . Note that $\langle\langle D[P_{\text{emp}}^{\xi}, P] \rangle\rangle_{N, \xi}$ data depends only on the set \mathcal{M} of prepared probability models. Since P_{GT} is unknown, and we estimate using only the sample $\{\xi_i\}_{i=1, \dots, N}$, it is natural that the ground truth P_{GT} does not explicitly affect $\langle\langle D[P_{\text{emp}}^{\xi}, P] \rangle\rangle_{N, \xi}$ data, which is irrelevant to estimation accuracy.

III. PERFORMANCE INEQUALITY

In the present section, we confirm that $\beta = N$ gives the best estimation in our framework, in which an average model weighted by KL divergence provides the estimation. In other words, we show that the expectation of the KL divergence between the grand truth P_{GT} and the estimated density \hat{P}_{β}^{ξ} , i.e., $\langle D[P_{\text{GT}}, \hat{P}_{\beta}^{\xi}] \rangle_{\text{data}}$ becomes minimum at $\beta = N$, where the estimation \hat{P}_{β}^{ξ} is an averaged model weighted by the KL divergence to the empirical distribution given by the data $\{\xi_i\}_{i=1, \dots, N}$ with the inverse temperature β :

$$\hat{P}_{\beta}^{\xi}(x) = \frac{\sum_{P \in \mathcal{M}} P(x) e^{-\beta D[P_{\text{emp}}^{\xi}, P]}}{\sum_{P \in \mathcal{M}} e^{-\beta D[P_{\text{emp}}^{\xi}, P]}}. \quad (20)$$

Considering the difference between $\langle D[P_{\text{GT}}, \hat{P}_{\beta}^{\xi}] \rangle_{\text{data}}$ and $\langle D[P_{\text{GT}}, \hat{P}_N^{\xi}] \rangle_{\text{data}}$, we have

$$\begin{aligned} & \langle D[P_{\text{GT}}, \hat{P}_{\beta}^{\xi}] \rangle_{\text{data}} - \langle D[P_{\text{GT}}, \hat{P}_N^{\xi}] \rangle_{\text{data}} \\ &= - \int \prod_{i=1}^N [d\xi_i P_{\text{GT}}(\xi_i)] \int dx P_{\text{GT}}(x) \ln \frac{\hat{P}_{\beta}^{\xi}(x)}{\hat{P}_N^{\xi}(x)}. \end{aligned} \quad (21)$$

Note that the ground truth changes as

$$\prod_{i=1}^N [d\xi_i P_{\text{GT}}(\xi_i)] P_{\text{GT}}(x) \rightarrow \prod_{i=1}^N d\xi_i P_{\text{GT}}^f(x) E_{P_{\text{GT}}^f}^{\xi, N} \quad (22)$$

under the gauge transformation, while $\hat{P}_{\beta}^{\xi} / \hat{P}_N^{\xi}$ is invariant. Since the result in Eq. (21) does not change before and after the gauge transformation, it equals the average over all gauges $f \in \mathcal{F}$. Since \mathcal{F} is the set of bijective maps connecting P_{GT} and $P \in \mathcal{M}$ as $P_{\text{GT}}^f(x) = P(x)$ in our framework, the average over $f \in \mathcal{F}$ equals the average over $P \in \mathcal{M}$. Then we find

$$\begin{aligned} & - \int \prod_{i=1}^N [d\xi_i P_{\text{GT}}(\xi_i)] \int dx P_{\text{GT}}(x) \ln \frac{\hat{P}_{\beta}^{\xi}(x)}{\hat{P}_N^{\xi}(x)} \\ &= - \frac{1}{N_{\mathcal{M}}} \int \prod_{i=1}^N d\xi_i \int dx S_{\mathcal{M}}^{\xi, N}(x) \ln \frac{\hat{P}_{\beta}^{\xi}(x)}{\hat{P}_N^{\xi}(x)} \\ &\geq - \frac{1}{N_{\mathcal{M}}} \int \prod_{i=1}^N d\xi_i \int dx S_{\mathcal{M}}^{\xi, N}(x) \left[\frac{\hat{P}_{\beta}^{\xi}(x)}{\hat{P}_N^{\xi}(x)} - 1 \right], \end{aligned} \quad (23)$$

where $S_{\mathcal{M}}^{\xi, \beta}(x) := \sum_{P \in \mathcal{M}} P(x) E_P^{\xi, \beta}$, and $\ln z \geq z - 1$ for $z > 0$ has been used.

Using the definition of the KL divergence $D[P_{\text{emp}}^{\xi}, P]$, the concrete form of $\hat{P}_{\beta}^{\xi}(x)$ is rewritten as

$$\hat{P}_{\beta}^{\xi}(x) = \frac{S_{\mathcal{M}}^{\xi, \beta}(x)}{\sum_{P \in \mathcal{M}} E_P^{\xi, \beta}}. \quad (24)$$

Thus, we evaluate the last line in Eq. (23) as

$$- \frac{1}{N_{\mathcal{M}}} \int \prod_{i=1}^N d\xi_i \int dx \left[S_{\mathcal{M}}^{\xi, \beta}(x) \frac{\sum_{P \in \mathcal{M}} E_P^{\xi, N}}{\sum_{P \in \mathcal{M}} E_P^{\xi, \beta}} - S_{\mathcal{M}}^{\xi, N}(x) \right]. \quad (25)$$

Since $\int dx P(x) = 1$ holds for an arbitrary model P , we have

$$\int dx S_{\mathcal{M}}^{\xi, \beta}(x) = \sum_{P \in \mathcal{M}} E_P^{\xi, \beta} \quad (26)$$

for arbitrary β . Using this fact, we find that Eq. (25) vanishes. Therefore, we attain

$$\langle D[P_{\text{GT}}, \hat{P}_{\beta}^{\xi}] \rangle_{\text{data}} \geq \langle D[P_{\text{GT}}, \hat{P}_N^{\xi}] \rangle_{\text{data}}. \quad (27)$$

The density function \hat{P}_N^{ξ} estimated on the Nishimori condition is closest to the ground truth P_{GT} in the sense of expectation.

The above result is suggestive. The most accurate estimation in the limit of $N \rightarrow \infty$ corresponds to the Nishimori temperature $\beta \rightarrow \infty$, or zero temperature estimation. It is consistent with the well-known fact that the best model is obtained through the conventional maximum likelihood method when the sample size is sufficiently large. In such a case, the criterion for model selection is minimization for the KL divergence to the empirical distribution P_{emp}^{ξ} determined from data $\{\xi_i\}_{i=1, \dots, N}$. In the conventional maximum likelihood method, a new model P with lower KL divergence replaces an old one with higher KL divergence. On the other hand, in our approach, we weigh these old models by the KL divergence to obtain a more accurate model.

To reduce the value of $\langle D[P_{\text{GT}}, \hat{P}_N^{\xi}] \rangle_{\text{data}}$, we should increase the sample size N or improve the set of models \mathcal{M} under consideration. Note that \hat{P}_N^{ξ} does not depend on P_{GT} but on the set of models \mathcal{M} . In other words, the same \mathcal{M} yields the same estimation result \hat{P}_N^{ξ} , even if the same sample $\{\xi_i\}_{i=1, \dots, N}$ is generated from different ground truths. It is a natural result because our estimation is based only on the data. The choice of \mathcal{M} determines the estimation accuracy. In Sec. VI, we discuss the improvement of \mathcal{M} . $\langle D[P_{\text{GT}}, \hat{P}_N^{\xi}] \rangle_{\text{data}}$ is an indicator for the estimation accuracy and depends on P_{GT} , although \hat{P}_N^{ξ} is independent of P_{GT} . We have shown that $\langle D[P_{\text{GT}}, \hat{P}_{\beta}^{\xi}] \rangle_{\text{data}}$ gives the minimum at $\beta = N$ using gauge invariance for the difference in Eq. (21). However, since $D[P_{\text{GT}}, \hat{P}_{\beta}^{\xi}]$ is not gauge invariant, the value of $\langle D[P_{\text{GT}}, \hat{P}_{\beta}^{\xi}] \rangle_{\text{data}}$ depends on P_{GT} .

IV. TWO POSSIBLE INTERPRETATIONS

We show in Sec. V that the estimation \hat{P}_N^{ξ} by our method is generally not in the prepared set \mathcal{M} , i.e., $\hat{P}_N^{\xi} \notin \mathcal{M}$. In this case, the best model does not belong to the initially prepared set of models. On the other hand, the conventional maximum likelihood method selects the model most explanatory to the realized data from the initially prepared models.

Thus, our framework requires a change in the interpretation of maximum likelihood estimation. There are two possible interpretations of our method. The first is to select the model probabilistically. In our approach, the prepared models have Boltzmann weights. We can select a model probabilistically according to the Boltzmann weight if only one of the models in \mathcal{M} is allowed as an estimation result. In this case, a lower KL divergence is achieved compared to one obtained by the conventional maximum likelihood method in the sense of expectation. The second is to accept $\hat{P}_N^\xi \notin \mathcal{M}$ as a probability model. In other words, we construct a new probability model from the set of initially prepared models $P \in \mathcal{M}$. In this case, the best estimation is deterministic. If we have prior knowledge on P_{GT} , the initial set of models \mathcal{M} is naturally prepared based on prior knowledge. However, the newly constructed model \hat{P}_N^ξ does not generally belong to the set of models consistent with the prior knowledge. We discuss such an example in Sec. V.

V. SIMPLE EXAMPLES

Based on the discussion up to Sec. III, the best estimation for the given data $\{\xi_i\}_{i=1,\dots,N}$ in our framework $\hat{P}_N^\xi(x)$ is given as

$$\hat{P}_N^\xi(x) = \frac{\sum_P P(x) \prod_{i=1}^N P(\xi_i)}{\sum_P \prod_{i=1}^N P(\xi_i)}. \quad (28)$$

In the present section, we deal with simple examples in which \hat{P}_N^ξ is tractable and we discuss the implications of our method.

A. Model selection from two

Let us estimate the probability density behind the given one-dimensional data $\{\xi_i\}_{i=1,\dots,N}$. We initially prepare two one-dimensional normal distributions $P_\pm(x) := \exp[-(x \mp a)^2/2]/\sqrt{2\pi}$ ($a > 0$). In other words, we set $\mathcal{M} = \{P_+, P_-\}$. In this case, simple calculation yields

$$\hat{P}_N^\xi(x) = \frac{e^{-\frac{1}{2}(x^2+a^2)}}{\sqrt{2\pi}} [\cosh ax + (\sinh ax)(\tanh aN\bar{\xi})], \quad (29)$$

where $\bar{\xi} := \frac{1}{N} \sum_{i=1}^N \xi_i$ is the sample average. Since $\tanh aN\bar{\xi} \rightarrow \text{sgn } \bar{\xi}$ in the limit of $N \rightarrow \infty$ with $a > 0$, we find $\hat{P}_N^\xi(x) \rightarrow \exp[-\frac{1}{2}(x - a \text{sgn } \bar{\xi})^2]/\sqrt{2\pi}$ in the large- N limit. Thus, our method selects one from the two models depending on $\bar{\xi}$ without uncertainty for $N \rightarrow \infty$.

As shown in Fig. 1, $\hat{P}_N^\xi(x)$ differs from the normal distribution for small N and converges rapidly to the normal distribution as N increases. The deviation from the normal distribution for small N can be interpreted as an interference between the two models. Since we cannot select one model definitely for a sample with a finite size, additional variance emerges due to the other model.

In addition to the effect of finite N , it is interesting that \hat{P}_N^ξ is given even for $\bar{\xi} = 0$. In the conventional maximum likelihood method, the KL divergence from the empirical distribution to P_+ equals the KL divergence to P_- in such a situation. Then one appropriate model cannot be selected as the solution to the optimization problem. On the other hand, a




FIG. 1. $\hat{P}_N^\xi(x)$ obtained for a set of two normal distributions. The results for $N = 2$, $N = 10$, and $N = 50$ are represented by curves formed with circles, triangles, and squares, respectively. The solid curve depicts the normal distribution $P_+(x) = \exp[-\frac{1}{2}(x - a)^2]/\sqrt{2\pi}$. The curves for \hat{P}_N^ξ with $N = 50$ and P_+ are almost overlapped. The parameters are set to be $\bar{\xi} = 0.1$ and $a = 1$.

unique optimum solution becomes $\hat{P}_N^\xi(x) = \frac{1}{2}[P_+(x) + P_-(x)]$ in our framework. Although the estimated model does not belong to the initially prepared set of probability models, i.e., $\hat{P}_N^\xi \notin \mathcal{M} = \{P_+, P_-\}$, it is considered as an intuitively appropriate solution to the optimization problem.

B. Estimation from data generated by a normal distribution

In the present subsection, we assume the case where the ground truth P_{GT} is known to be a one-dimensional normal distribution. In such a situation, the natural set of models \mathcal{M} consists of all possible one-dimensional normal distributions, i.e., $P(x|\sigma, a) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp[-\frac{1}{2\sigma^2}(x - a)^2]$ defined for all positive numbers $\sigma > 0$ and real numbers $a \in \mathbb{R}$.

Defining the set of probability models by $\mathcal{M} := \{P(x|\sigma, a) | \sigma > 0, a \in \mathbb{R}\}$, we evaluate the denominator in Eq. (28) as

$$Z_{\mathcal{M}}^\xi = \frac{1}{2} C_N U_N(V_\xi) \Gamma\left(\frac{N}{2} - 1\right), \quad (30)$$

where $V_\xi := \frac{1}{N} \sum_{i=1}^N \xi_i^2 - \bar{\xi}^2$ is the sample variance for the data ξ , and $\Gamma(z) := \int_0^\infty dt t^{z-1} e^{-t}$ is a Gamma function. We have used the notation $Z_{\mathcal{M}}^\xi := \sum_{P \in \mathcal{M}} \prod_{i=1}^N P(\xi_i)$. In addition, the notations $C_N := [(2\pi)^{(N-1)/2} \sqrt{N}]^{-1}$ and $U_N(z) := [2/(Nz)]^{N/2-1}$ have been used. Similarly, the numerator in Eq. (28) is

$$S_{\mathcal{M}}^\xi = \frac{1}{2} C_{N+1} U_{N+1}(V_{\xi,x}) \Gamma\left(\frac{N}{2} - \frac{1}{2}\right), \quad (31)$$

where $V_{\xi,x} := \frac{1}{N+1} (N\bar{\xi} + x)^2 - [\frac{1}{N+1} (N\bar{\xi} + x)]^2$ is the sample variance for the data composed of ξ and the new sample point x . The notation $S_{\mathcal{M}}^\xi(x) := S_{\mathcal{M}}^{\xi,N}$ has also been used. Then the




FIG. 2. $\hat{P}_N^\xi(x)$ for \mathcal{M} consisting of all normal distributions. The results for $N = 3$, $N = 10$, and $N = 50$ are represented by curves formed with circles, triangles, and squares, respectively. The solid curve depicts the normal distribution $\exp(-\frac{1}{2}x^2)/\sqrt{2\pi}$. The parameters are set to be $\bar{\xi} = 0$ and $V_\xi = 1$.

best estimation on the Nishimori condition becomes

$$\hat{P}_N^\xi(x) = \frac{C_{N+1}U_{N+1}(V_{\xi,x})}{C_N U_N(V_\xi)} \frac{\Gamma(N/2 - 1/2)}{\Gamma(N/2 - 1)}. \quad (32)$$

Since

$$\hat{P}_N^\xi(x) \propto \left[\frac{(N+1)V_\xi}{(x-\bar{\xi})^2 + (N+1)V_\xi} \right]^{(N-1)/2}, \quad (33)$$

it is found that \hat{P}_N^ξ is given by $(N-1)/2$ power of the Lorentzian distribution with the peak located at $\bar{\xi}$ and the peak width of $[(N+1)V_\xi]^{1/2}$ as shown in Fig. 2. The obtained model is not a normal distribution, although our estimation assumes that the ground truth P_{GT} obeys normal distribution. This result seems to be counterintuitive. However, the prior knowledge of the ground truth of normal distribution does not

limit the model to normal distribution. When the sample size N is finite, selecting one normal distribution may lead to a large deviation from the ground truth P_{GT} . To rule out the possibility of a large deviation, the best estimation deviates from any normal distributions. We emphasize that such a situation occurs because of the finite sample size N . Indeed, we find in the limit $N \rightarrow \infty$ that

$$\hat{P}_N^\xi(x) \rightarrow \frac{1}{\sqrt{2\pi V_\xi}} \exp\left[-\frac{1}{2V_\xi}(x-\bar{\xi})^2\right] \quad (34)$$

holds using Stirling's formula $\Gamma(z+1) \sim \sqrt{2\pi z}(z/e)^z$ and the definition of the exponential function $(1+z/n)^n \rightarrow e^z$ for $n \rightarrow \infty$. Thus, in the limit of large samples, our result reproduces the normal distribution, which the prior knowledge implicates, and is consistent with the well-known conclusion of the maximum likelihood estimation. \hat{P}_N^ξ given by Eq. (32) can be interpreted as a finite size effect version of normal distribution.

VI. IMPROVEMENT OF MODEL SET

The best estimation \hat{P}_N^ξ depends only on the set of prepared models \mathcal{M} . Let us consider the sequential extension of the set of models \mathcal{M} to improve the estimation accuracy. In this process, a new set of models \mathcal{M}_* is constructed by the addition of a new model P_* to the original set \mathcal{M} . If the number $N_{\mathcal{M}}$ of elements in \mathcal{M} is sufficiently large, the constructed model $\hat{P}_{N,*}^\xi$ from the modified set \mathcal{M}_* is expressed as

$$\hat{P}_{N,*}^\xi(x) = \hat{P}_N^\xi(x) + [P_*(x) - \hat{P}_N^\xi(x)] \frac{z_*^\xi}{Z_{\mathcal{M}}^\xi} + \mathcal{O}((z_*^\xi/Z_{\mathcal{M}}^\xi)^2), \quad (35)$$

where $z_*^\xi := \prod_{i=1}^N P_*(\xi_i)$. For practical purposes, it is preferable $\hat{P}_{N,*}^\xi$ is closer to the ground truth P_{GT} than \hat{P}_N^ξ is in the KL divergence. The expectation of the KL divergence between the newly obtained model $\hat{P}_{N,*}^\xi$ and P_{GT} is calculated as

$$\langle D[P_{\text{GT}}, \hat{P}_{N,*}^\xi] \rangle_{\text{data}} = \langle D[P_{\text{GT}}, \hat{P}_N^\xi] \rangle_{\text{data}} - \int \prod_{i=1}^N [d\xi_i P_{\text{GT}}(\xi_i)] \int dx [P_*(x) - \hat{P}_N^\xi(x)] \frac{P_{\text{GT}}(x) z_*^\xi}{\hat{P}_N^\xi(x) Z_{\mathcal{M}}^\xi} + \mathcal{O}((z_*^\xi/Z_{\mathcal{M}}^\xi)^2). \quad (36)$$

Since $\hat{P}_N^\xi(x)$ is non-negative, it is necessary to add P_* such that $P_*(x) - \hat{P}_N^\xi(x)$ becomes as large as possible at x where $P_{\text{GT}}(x)/\hat{P}_N^\xi(x)$ is large to make $\langle D[P_{\text{GT}}, \hat{P}_{N,*}^\xi] \rangle_{\text{data}}$ as small as possible. In other words, it is preferable to add a model that assigns as high a probability as possible at x , which is underestimated in the originally constructed model \hat{P}_N^ξ . The systematic expansion of \mathcal{M} described above leads to a decrease in the expected value of the KL divergence between \hat{P}_N^ξ and P_{GT} . Such a systematic process may yield an algorithm that can efficiently estimate P_{GT} .

VII. RELATION TO BAYESIAN UPDATING

Our proposed method is related to Bayesian updating. It updates the prior distribution $\hat{p}(\theta)$ for the parameter θ of the model $P_\theta(x)$ based on observed data:

$$\hat{p}(\theta|\xi) = \frac{P_\theta(\xi)\hat{p}(\theta)}{\sum_\theta P_\theta(\xi)\hat{p}(\theta)}. \quad (37)$$

In our method, the weights for probability models are updated when new data ξ_{N+1} are observed. According to our result,

Eq. (28), the weight for the model P is given as

$$\hat{p}(P|\xi) = \frac{\prod_{i=1}^N P(\xi_i)}{\sum_P \prod_{i=1}^N P(\xi_i)} \quad (38)$$

when the realization of the data $\{\xi_i\}_{i=1,\dots,N}$ is observed. Interpreting $\hat{P}(P|\xi_1, \dots, \xi_N)$ as the posterior distribution of the model P , when new data ξ_{N+1} are added, the posterior distribution of the model P is given by Bayesian updating as

$$\hat{p}(P|\xi_1, \dots, \xi_N, \xi_{N+1}) = \frac{P(\xi_{N+1})\hat{p}(P|\xi_1, \dots, \xi_N)}{\sum_P P(\xi_{N+1})\hat{p}(P|\xi_1, \dots, \xi_N)}. \quad (39)$$

Note that the update rule, Eq. (39), is a direct result from Eq. (28). Our proposed method is concluded to be consistent with Bayesian updating starting from a prior distribution of equal weights for all probability models.

By repeating the Bayesian updating with given data $\{\xi_i\}_{i=1,\dots,N}$ fixed, we can improve the posterior distribution for the probability model P by obtaining the posterior distribution. In this approach, the posterior distribution for the model P of the $(k+1)$ th step $\hat{p}_{k+1}(P|\xi_1, \dots, \xi_N)$ is obtained from the posterior distribution of the k th step $\hat{p}_k(P|\xi_1, \dots, \xi_N)$ as the prior distribution. $\hat{p}_{k+1}(P|\xi_1, \dots, \xi_N)$ is updated as

$$\hat{p}_{k+1}(P|\xi_1, \dots, \xi_N) = \frac{\hat{p}_k(P|\xi_1, \dots, \xi_N) \prod_{i=1}^N P(\xi_i)}{\sum_P \hat{p}_k(P|\xi_1, \dots, \xi_N) \prod_{i=1}^N P(\xi_i)}, \quad (40)$$

where $\hat{p}_0(P|\xi_1, \dots, \xi_N) = \hat{p}_0(P)$ is the initial prior distribution for the model P . For this update rule, the fixed point $\hat{p}_\infty(P|\xi_1, \dots, \xi_N)$ is given by Eq. (38). Thus, our method can be understood as a fixed point of Bayesian updating.

It has been pointed out that Bayesian updating does not require regularization. We have found that Bayesian updating can be interpreted as an error-correcting version of the maximum likelihood method. In other words, we can conclude that Bayesian updating is an estimation method that minimizes the KL divergence between the ground truth distribution and the probability model in the sense of expected value.

VIII. CONCLUSION

We have discussed a method for estimating density functions based on gauge symmetry in the KL divergence. Our framework naturally prevents overfitting, although the conventional maximum likelihood method can lead to incorrect estimation results due to overfitting to empirical distributions. While empirical regularization to prevent overfitting often requires a search for the optimal hyperparameters, our method does not require adjustment of hyperparameters because the optimal hyperparameter is given explicitly as the Nishimori condition. As a natural extension of the conventional maximum likelihood method, our estimation \hat{P}_N^ξ is realized by the average of probability models with Boltzmann weights. The result \hat{P}_N^ξ is the best estimation obtained from the sample $\{\xi_i\}_{i=1,\dots,N}$ as shown in Sec. III. The performance of the estimations \hat{P}_N^ξ is determined by the set of prepared models \mathcal{M} . Therefore, it is necessary to find \mathcal{M} that provides high estimation accuracy. In Sec. VI, we have discussed the method

of sequential extension to obtain such a set. In addition, we have discussed two interpretations of our method in Sec. IV and simple examples in Sec. V.

Our method is valid when the data $\{\xi_i\}_{i=1,\dots,N}$ are given independently. On the other hand, it cannot be applied directly to a sequence of correlated data encountered in many practical situations. In the sense of prediction, there is no use in estimating the distribution $P_{\text{GT}}(\xi_i)$ for a single datum ξ_i when $\{\xi_i\}_{i=1,\dots,N}$ are correlated. For a sequence of correlated data, it would be more meaningful to estimate the distribution $P_{\text{GT}}(\xi_1, \dots, \xi_N)$ of the series of the entire data series $\{\xi_i\}_{i=1,\dots,N}$. In this case, we should focus on estimating the mechanism that generates the data series. For example, for the series of data $\{\xi_i\}_{i=1,\dots,N}$ generated by a stationary Markov process, estimating the transition probability $T(\xi_{i+1}|\xi_i)$ is more fundamental than estimating the distribution $P_{\text{GT}}(\xi_1, \dots, \xi_N)$. In this case, ξ_{i+1} can be considered to be generated independently with a given state ξ_i . Thus, our method can be applied to estimate the transition probability $T(\xi_{i+1}|\xi_i)$.

Through our proposed method, the average of multiple models weighted by KL divergence gives a better model. However, it is impractical to calculate \hat{P}_N^ξ concretely, in general, because a partition function differs for each model. For example, even if the models are limited to the Boltzmann machines [27,28], it is not easy to calculate a partition function for each model. On the other hand, such an average of multiple models is often used in ensemble learning [29]. Therefore, the difficulties faced in model averaging are the same as those encountered in ensemble learning. Our method is practically applicable for a set of probability models that are handled in ensemble learning. Ensemble learning tries to overcome them with various ideas. For example, parallel computation is effective since each model can be built independently. Random forests take a simple arithmetic mean of the expected values from many models, but from our point of view, we conclude that an average weighted by KL divergence is more reliable. Our proposed method is expected to help improve the accuracy of ensemble learning.

The gauge symmetry of the KL divergence plays a central role in the maximum likelihood estimation. In supervised machine learning with a loss function, it would be possible to perform a regularization similar to our method by exploiting the gauge symmetry of the loss function. In the training phase of conventional machine learning with many parameters, such as deep learning, the model that minimizes the loss function is solely selected. In such a case, any loss function (and any regularization) gives almost the same results as long as the model is not significantly different from the ground truth. A model with high-generalization performance is obtained empirically with little dependence on the form of the loss function (and regularization term) in deep learning with sufficiently large \mathcal{M} [30,31]. However, in our framework, the generalization performance is expected to be highly dependent on the form of the loss function since models with high loss also contribute to the results. It is future work to understand the relationship between the choice of the loss function and generalization performance in machine learning from the viewpoint of gauge symmetry.

- [1] S. Kullback and R. A. Leibler, *Ann. Math. Stat.* **22**, 79 (1951).
- [2] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, Berlin, 2006).
- [3] G. E. Hinton and D. Van Camp, in *Proceedings of the Sixth Annual Conference on Computational Learning Theory* (ACM, New York, 1993), pp. 5–13.
- [4] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [5] R. Tibshirani, *J. R. Stat. Soc.: Ser. B* **58**, 267 (1996).
- [6] A. E. Hoerl and R. W. Kennard, *Am. J. Math. Manag. Sci.* **1**, 5 (1981).
- [7] H. Zou and T. Hastie, *J. R. Stat. Soc.: Ser. B* **67**, 301 (2005).
- [8] C.-H. Zhang, *Ann. Stat.* **38**, 894 (2010).
- [9] H. Akaike, *IEEE Trans. Autom. Control* **19**, 716 (1974).
- [10] K. P. Burnham and D. R. Anderson, *Sociol. Methods Res.* **33**, 261 (2004).
- [11] S. Konishi and G. Kitagawa, *Information Criteria and Statistical Modeling*, 1st ed. (Springer, Berlin, 2007).
- [12] R. W. Hamming, *Bell Syst. Tech. J.* **29**, 147 (1950).
- [13] I. S. Reed and G. Solomon, *J. Soc. Ind. Appl. Math.* **8**, 300 (1960).
- [14] R. G. Gallager, *Information Theory and Reliable Communication* (Wiley & Sons, New York, 1968).
- [15] D. Saad, Y. Kabashima, T. Murayama, and R. Vicente, in *Cryptography and Coding* (Springer, Berlin, 2001), pp. 307–316.
- [16] M. Bayati and A. Montanari, *IEEE Trans. Inf. Theory* **57**, 764 (2011).
- [17] N. Surlas, *Nature (London)* **339**, 693 (1989).
- [18] P. Ruján, *Phys. Rev. Lett.* **70**, 2968 (1993).
- [19] H. Nishimori and K. Y. Michael Wong, *Phys. Rev. E* **60**, 132 (1999).
- [20] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, Oxford, 2001).
- [21] H. Nishimori, *Prog. Theor. Phys.* **66**, 1169 (1981).
- [22] P. Le Doussal and A. B. Harris, *Phys. Rev. Lett.* **61**, 625 (1988).
- [23] P. Contucci, C. Giardinà, and H. Nishimori, in *Spin Glasses: Statics and Dynamics*, edited by A. B. de Monvel and A. Bovier (Birkhäuser, Basel, 2009), pp. 103–121.
- [24] T. Sei, *Ann. Inst. Stat. Math.* **63**, 675 (2011).
- [25] Y. Brenier, *Commun. Pure Appl. Math.* **44**, 375 (1991).
- [26] R. J. McCann, *Duke Math. J.* **80**, 309 (1995).
- [27] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cognit. Sci.* **9**, 147 (1985).
- [28] G. E. Hinton, in *Neural Networks: Tricks of the Trade*, edited by G. Montavon, G. B. Orr, and K. R. Müller (Springer, Berlin, 2012).
- [29] C. Zhang and Y. Ma, *Ensemble Machine Learning: Methods and Applications* (Springer, Berlin, 2012).
- [30] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, *Commun. ACM* **64**, 107 (2021).
- [31] C. Zhang, Q. Liao, A. Rakhlin, K. Sridharan, B. Miranda, N. Golowich, and T. Poggio, Memo No. 067, December 26, 2017, Musings on deep learning: Properties of SGD, Technical Report, Center for Brains, Minds and Machines (CBMM), 2017.