

---

# Gibbs-based Information Criteria and the Over-Parameterized Regime

---

Haobo Chen<sup>†</sup>

<sup>†</sup> University of Florida

Gregory W. Wornell<sup>‡</sup>

<sup>‡</sup> Massachusetts Institute of Technology

Yuheng Bu<sup>†</sup>

## Abstract

Double-descent refers to the unexpected drop in test loss of a learning algorithm beyond an interpolating threshold with over-parameterization, which is not predicted by information criteria in their classical forms due to the limitations in the standard asymptotic approach. We update these analyses using the information risk minimization framework and provide Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) for models learned by the Gibbs algorithm. Notably, the penalty terms for the Gibbs-based AIC and BIC correspond to specific information measures, i.e., symmetrized KL information and KL divergence. We extend this information-theoretic analysis to over-parameterized models by providing two different Gibbs-based BICs to compute the marginal likelihood of random feature models in the regime where the number of parameters  $p$  and the number of samples  $n$  tend to infinity, with  $p/n$  fixed. Our experiments demonstrate that the Gibbs-based BIC can select the high-dimensional model and reveal the mismatch between marginal likelihood and population risk in the over-parameterized regime, providing new insights to understand double-descent.

## 1 INTRODUCTION

The classical understanding of model selection is that more complex models can capture more complex patterns but tend to overfit and have large generalization errors (Geman et al., 1992). This tradeoff results in a U-shaped curve which is characterized by the classical model selection criterion when test loss is plotted

against model complexity. As a result, the models that minimize test loss tend to have moderate complexity. Recently, the success of deep learning challenges this classical picture since neural networks are often extremely complex (e.g., able to fit random labels (Zhang et al., 2017)) while *also* generalizing well to yield low test error on unseen samples.

An emerging explanation of this behavior is *double-descent* (Belkin et al., 2019), which posits that: 1) The classical U-shaped curve is only valid when the number of model parameters  $p$  is smaller than the number of samples  $n$ . 2) In the over-parameterized regime where  $p$  is significantly larger than  $n$ , and models are complex enough to fit training data perfectly, test loss can decrease with increased model complexity.

To better understand the double-descent phenomenon, we revisit the classical derivations of information criteria and discern that the penalty term in Akaike Information Criterion (AIC) can be interpreted as the generalization error within a broader learning context, while Bayesian Information Criterion (BIC) approximates the marginal likelihood using the empirical risk minimization solution. We further update the classical analyses of AIC and BIC using the information risk minimization framework proposed in Zhang (2006) by focusing on the optimal Gibbs algorithm (distribution). It is shown in Aminian et al. (2021) that the generalization error of the Gibbs algorithm can be characterized using information measures. This information-theoretic analysis motivates the proposed Gibbs-based information criteria, which readily extend to over-parameterized models.

We make the following contributions in this paper:

1. We provide information-theoretic analyses for the generalization error and marginal likelihood of the model learned by the Gibbs algorithm, resulting in Gibbs-based AIC (14) and Gibbs-based BIC<sup>+</sup> (17) and BIC<sup>-</sup> (18), with different information measures as the penalty terms.
2. We show that the Gibbs-based information criteria align with the classical information criteria in the classical large  $n$  regime theoretically (Theorem 1 and 2) and empirically.

---

Proceedings of the 27<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

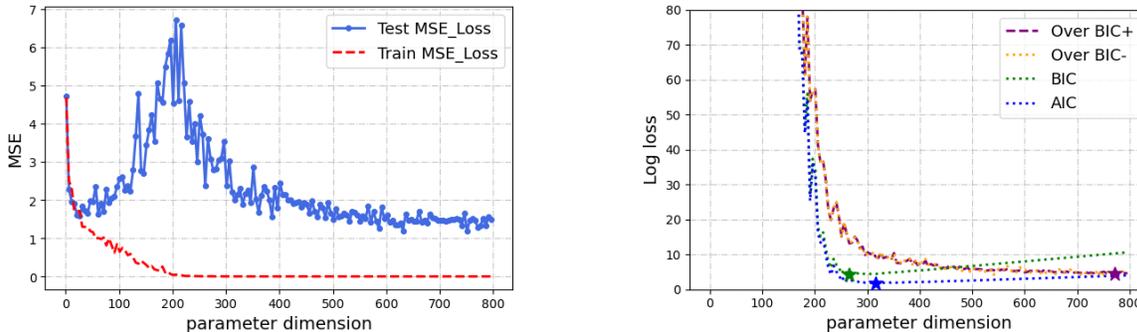


Figure 1: The test and training mean-squared error (MSE) (**left**) and the comparison of the proposed over-parameterized  $\text{BIC}^+$  and  $\text{BIC}^-$  with other classical information criteria (**right**) for over-parameterized RF models plotted with varying parameter dimension  $p$ . The preferred models selected by different information criteria are marked using stars with different colors. More details can be found in Section 5.

3. We generalize our information-theoretic analysis to over-parameterized random feature (RF) models, which results in over-parameterized Gibbs-based BICs ((28) and (29)) that favor over-parameterized RF model, while classical information criteria cannot; see Figure 1.
4. We empirically compare the Gibbs-based BICs and AIC in the over-parameterized RF model by decomposing them into different terms (Section 5), and show the mismatch between marginal likelihood (BIC) and generalization error (AIC) in the over-parameterized setting, where AIC exhibits double-descent but BIC does not. Such a phenomenon is highly affected by the choice of prior distributions.

**Related Work** Previous work has extended the classical BIC to high dimensions, e.g., Chen and Chen (2008); Fan and Tang (2013). However, these works seek to *replace* maximizing marginal likelihood with a different criterion, substituting a penalty term  $pf(n)$  in place of the  $p \log n$  in BIC. A Widely Applicable BIC (WBIC) is proposed in Watanabe (2013), which extends the BIC to singular probabilistic models using the asymptotic approximation of the marginal likelihood with Gibbs posterior in the classical large  $n$  regime. By contrast, we retain the BIC criterion but analyze it beyond the classical regime for the Gibbs algorithm with an exact information-theoretic analysis.

Double-descent of the population risk with increasing model size was introduced in Belkin et al. (2019); see also Advani et al. (2020); Geiger et al. (2019). An empirical demonstration of double-descent in modern deep networks is provided in Nakkiran et al. (2019). A variety of work develops simplified models where the characterization of the double-descent curve can be obtained. For example, double-descent in linear regression models is investigated in Belkin et al. (2020);

Hastie et al. (2022); Bartlett et al. (2020); Muthukumar et al. (2020), and in linear classification models in Deng et al. (2022); Kini and Thrampoulidis (2020); Gerace et al. (2020). The RF model has been adopted to understand double-descent in Mei and Montanari (2022), which provides a generalization analysis of the performance achieved with ridge regression in the over-parameterized regime. In some of our analyses, we likewise adopt this RF model (Mei and Montanari, 2022; d’Ascoli et al., 2020; Gerace et al., 2020; Weinan et al., 2020; Liu et al., 2022), but with a different objective and analysis tools.

Double-descent phenomena have been explained from different perspectives. In d’Ascoli et al. (2020); Yang et al. (2020); Dwivedi et al. (2020), double-descent curves are explained via a refined version of bias-variance tradeoff, where the bias of the model decreases monotonically with the increase of  $p$ , but the variance increases and then decreases with  $p$ . And the connection of gradient descent dynamics and double-descent is discussed in Weinan et al. (2020); Advani et al. (2020). In Nakkiran (2019); d’Ascoli et al. (2020), sample-wise double-descent is studied under linear regression, and Heckel and Yilmaz (2020) shows that by adjusting the step sizes, sample-wise double-descent can be eliminated by early stopping.

Moreover, Germain et al. (2016) connects the PAC-Bayes generalization bound with marginal likelihood and approximates the KL divergence for model selection using Markov chain Monte Carlo (MCMC) algorithms. However, they did not consider the over-parameterized regime. For over-parameterized models, Hodgkinson et al. (2023) study the Interpolating BIC with a special focus on interpolated models with zero training loss. Our paper is most related to Immer et al. (2021); Lotfi et al. (2022), which also examines the difference between marginal likelihood and generalization error

in model selection. However, by focusing on the Gibbs algorithm, we are able to interpret the mismatch between AIC and BIC via information measures, which is more insightful in understanding double descent and other complex behaviors.

## 2 PRELIMINARIES

We consider the following standard supervised learning formulation. Let  $S = \{Z_i\}_{i=1}^n$  be the training set, where each sample  $Z_i = \{(X_i, Y_i)\} \in \mathcal{Z}$  are i.i.d. generated from the data distribution  $P_Z$ , and with the realization of this sequence denoted as  $\mathbf{z}^n = (\mathbf{x}^n, y^n)$ . Denote the parameter of a machine learning model by  $w \in \mathcal{W}$ , where  $\mathcal{W}$  is the parameter space. The performance of the model is measured by a loss  $\ell: \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$ , and the log loss  $\ell_{\log}(w, \mathbf{z}) \triangleq -\log P(y|\mathbf{x}; w)$  associated with a parametric probabilistic model  $P(y|\mathbf{x}; w)$  is of particular interest to us.

We define the empirical risk and the population risk associated with a given  $w$  as

$$L_E(w, \mathbf{z}^n) \triangleq \frac{1}{n} \sum_{i=1}^n \ell(w, z_i), \quad (1)$$

$$L_P(w, P_Z) \triangleq \mathbb{E}_{P_Z}[\ell(w, Z)], \quad (2)$$

respectively. A learning algorithm can be modeled as a randomized mapping from the training set  $S$  onto a model parameter  $\hat{W}$  according to the conditional distribution  $P_{\hat{W}|S}$ .

Although the widely-used empirical risk minimization (ERM) is deterministic optimization, it is usually solved via stochastic gradient descent (SGD), which is random in nature. The ERM solution of log loss is the well-known maximum likelihood estimate (MLE), i.e.,  $\hat{W}_{\text{ML}} \triangleq \arg \max_w \hat{L}(w)$ , where  $\hat{L}(w) \triangleq \sum_{i=1}^n \log P_k(y_i|\mathbf{x}_i; w)$  denotes the log-likelihood of  $\mathbf{z}^n$ .

The expected generalization error quantifying the degree of over-fitting can be expressed in the form

$$\overline{gen}(P_{\hat{W}|S}, P_S) \triangleq \mathbb{E}_{P_{\hat{W}, S}}[L_P(\hat{W}, P_Z) - L_E(\hat{W}, S)], \quad (3)$$

where the expectation is taken over the joint distribution  $P_{\hat{W}, S} = P_{\hat{W}|S} \otimes P_S$ .

### 2.1 The Classical Forms of AIC and BIC

The standard derivation of the AIC and the BIC arises from the classical asymptotic analysis of MLE. Assume we have  $K$  candidate models  $M_1, M_2, \dots, M_K$ , and each model  $M_k$  is characterized by a parametric probabilistic model  $P_k(y|\mathbf{x}; \boldsymbol{\theta}_k)$ , a prior distribution  $\pi_k(\boldsymbol{\theta}_k)$ , where  $\boldsymbol{\theta}_k \in \mathcal{W}_k \subset \mathbb{R}^{p_k}$  is the parameter vector. We demonstrate that AIC selects the model with the

smallest population risk, and BIC identifies the true data-generating model by maximizing the marginal likelihood.

**AIC** This criterion (Akaike, 1981) ranks statistical models based on the Kullback-Leibler (KL) divergence between the true data distribution  $P_Z$  and the learned parametric model. With  $\hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)}$  denoting the MLE of the  $k$ th model, AIC selects the model as the solution to

$$\begin{aligned} & \arg \min_k D(P_Z \| P_k(y|\mathbf{x}; \hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)})) \\ & = \arg \min_k \mathbb{E}_{P_Z}[-\log P_k(y|\mathbf{x}; \hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)})]. \end{aligned} \quad (4)$$

The term  $\mathbb{E}_{P_Z}[-\log P_k(y|\mathbf{x}; \hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)})]$  can be interpreted as the population risk  $L_P(\hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)}, P_Z)$  of the MLE under log-loss. From this perspective, AIC measures how well the model fits the unknown data distribution  $P_Z$ , with smaller AIC values suggesting a lower population risk.

As the true distribution  $P_Z$  is unknown, the AIC is obtained by approximating the population risk as the sum of empirical risk, i.e., the negative log-likelihood of  $\hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)}$  on training samples and a penalty term corresponding to the generalization error. In the classic regime where  $p_k$  is fixed and  $n \rightarrow \infty$ , the asymptotic normality of MLE yields

$$\text{AIC} = -\frac{\hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}})}{n} + \frac{p}{n}. \quad (5)$$

Note that our form of AIC differs by a factor of 2 from its classical form to facilitate a direct comparison to population risk and generalization error as defined in (3). Detailed derivations are provided, for reference, in Appendix A.

**BIC** This criterion (Schwarz et al., 1978) ranks statistical models by their marginal likelihoods of generating the data, where smaller values of the BIC correspond to larger marginal likelihoods. Approximating the marginal likelihood of observing  $\mathbf{z}^n$  for  $M_k$ , i.e.,  $m_k(\mathbf{z}^n) \triangleq \int P_k(y^n|\mathbf{x}^n; \boldsymbol{\theta}_k) \pi_k(\boldsymbol{\theta}_k) d\boldsymbol{\theta}_k$ , Laplace's method yields

$$\log m_k(\mathbf{z}^n) = \hat{L}_k(\hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)}) - \frac{p_k}{2} \log n + O(1), \quad n \rightarrow \infty.$$

In turn, BIC is obtained by dropping terms that do not scale with  $n$  and scaling by  $-1/n$ :

$$\text{BIC} = -\frac{\hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}})}{n} + \frac{p \log n}{2n}. \quad (6)$$

When, further,  $P(M_k) = 1/K$ , we obtain

$$P(M_k|\mathbf{z}^n) = \frac{m_k(\mathbf{z}^n)P(M_k)}{\sum_{k=1}^K m_k(\mathbf{z}^n)P(M_k)} \propto m_k(\mathbf{z}^n). \quad (7)$$

Thus, when we assume a uniform prior over different models, the BIC ranks models by their posterior probability of generating the training data, and choosing the smallest BIC corresponds to the maximum a posteriori rule for model selection. A detailed derivation is provided in Appendix B.

Both (5) and (6) share a common first term, representing the average negative log-likelihood of the training data for MLE, which can be interpreted as the empirical risk with log-loss, decreasing as we adopt more complex models. We note that AIC and BIC in the classical  $n \rightarrow \infty$  regime are independent of the form of the model family  $P(y|\mathbf{x}; \boldsymbol{\theta})$  and the prior distribution  $\pi(\boldsymbol{\theta})$ , which makes it compatible with general distribution families subject to mild smoothness constraints. However, they select different models because AIC and BIC differ in the second penalty term.

## 2.2 Information Risk Minimization and Gibbs Algorithm

Classical AIC and BIC depend on MLE, which can be viewed as an ERM solution that purely minimizes empirical risk. Instead, we motivate the Gibbs algorithm using an information risk minimization framework, which minimizes both empirical risk and a generalization error bound.

We start with the following mutual information-based generalization error bound proposed in Xu and Raginsky (2017).

**Lemma 1.** (Xu and Raginsky, 2017) *Suppose the loss function  $\ell(w, z) \in [0, 1]$  is bounded, and  $S = \{Z_i\}_{i=1}^n$  contains  $n$  i.i.d. training samples, then  $|\overline{gen}(P_{\hat{W}|S}, P_S)| \leq \sqrt{I(\hat{W}; S)/(2n)}$ .*

From Lemma 1, we know that the mutual information between training data  $S$  and the learned parameter  $\hat{W}$  can be used as an upper bound for generalization error. Thus, Xu and Raginsky (2017) further considers the following algorithm that minimizes empirical risks regularized with mutual information,

$$P_{\hat{W}|S}^* = \operatorname{argmin}_{P_{\hat{W}|S}} \mathbb{E}_{P_{\hat{W}, S}} [L_E(\hat{W}, S)] + \frac{1}{\beta} I(\hat{W}; S), \quad (8)$$

where  $\beta > 0$  controls the regularization term and balances between over-fitting and generalization. As  $\beta \rightarrow \infty$ , it reduces back to the standard ERM algorithm.

As computing  $I(\hat{W}; S)$  requires the knowledge of  $P_{\hat{W}|S}$ , Zhang (2006); Xu and Raginsky (2017); Perlaza et al. (2022) further relax (8) by replacing it with an upper bound  $D(P_{\hat{W}|S} \| \pi | P_S) \geq I(\hat{W}; S)$ , where  $\pi$  is an arbitrary prior distribution over  $\mathcal{W}$ . The following lemma characterizes the solution to the relaxed problem.

**Lemma 2.** (Zhang, 2006; Perlaza et al., 2022) *The minimum value of the following information risk minimization (IRM) problem is*

$$\begin{aligned} \min_{P_{\hat{W}|S}} \mathbb{E}_{P_{\hat{W}, S}} [L_E(\hat{W}, S)] + \frac{1}{\beta} D(P_{\hat{W}|S} \| \pi | P_S) \\ = -\frac{1}{\beta} \mathbb{E}_{P_S} [\log \mathbb{E}_{\pi} [e^{-\beta L_E(W, S)}]], \end{aligned} \quad (9)$$

which is achieved by the following Gibbs algorithm

$$P_{\hat{W}|S}^*(w|s) = \frac{\pi(w)e^{-\beta L_E(w, s)}}{\mathbb{E}_{\pi} [e^{-\beta L_E(W, s)}]}, \quad \text{for } \beta > 0. \quad (10)$$

In addition, Perlaza et al. (2023) shows that the minimum value of IRM can be decomposed into the empirical risk under the prior distribution and the relative entropy  $D(\pi \| P_{\hat{W}|S})$ .

**Lemma 3.** (Perlaza et al., 2023) *The Gibbs distribution  $P_{\hat{W}|S}^*(w|s)$  in form (10) satisfies*

$$\begin{aligned} \mathbb{E}_{\pi \otimes P_S} [L_E(W, S)] - \frac{1}{\beta} D(\pi \| P_{\hat{W}|S}^* | P_S), \\ = -\frac{1}{\beta} \mathbb{E}_{P_S} [\log \mathbb{E}_{\pi} [e^{-\beta L_E(W, S)}]]. \end{aligned} \quad (11)$$

## 2.3 Sampling Algorithms for Gibbs Distribution

The Gibbs distribution was first proposed by Gibbs (1902) in statistical mechanics and further explored in the context of information theory by Jaynes (1957). In general, it is difficult to compute Gibbs posterior directly due to the integral in the partition function  $\mathbb{E}_{\pi} [e^{-\beta L_E(W, s)}]$ . In practice, we delve into two stochastic algorithms, stochastic gradient Langevin dynamics (SGLD) or Metropolis-adjusted Langevin algorithm (MALA), to sample from the Gibbs distribution.

**SGLD** The SGLD algorithm is defined using the following update formula,

$$\hat{W}_{t+1} = \hat{W}_t - \eta \nabla L_E(\hat{W}_t, s) + \sqrt{\frac{2\eta}{\beta}} \zeta_t, \quad t = 0, 1, \dots \quad (12)$$

where  $\zeta_t$  is a standard Gaussian random vector and  $\eta \geq 0$  is the step size. SGLD can also be viewed as a noisy version of standard SGD. It is shown in Raginsky et al. (2017) that under some conditions on loss function, the Wasserstein distance between the distribution  $P_{\hat{W}_t|S}$  induced by SGLD and the Gibbs distribution  $P_{\hat{W}|S}^*$  will converge to zero as  $t$  goes infinity. This allows us to sample  $\hat{W}$  from the Gibbs distribution using SGLD with a sufficiently large number of iterations  $t$ . Note that SGLD has been applied similarly to optimize generalization bound in Dziugaite and Roy (2018).

**MALA** Another approach is the Metropolis-adjusted Langevin algorithm (MALA) (Dwivedi et al., 2018). MALA and SGLD are first-order sampling methods since they have similar gradient update formulas, which guarantees that both algorithms converge to the Gibbs distribution. MALA differs from the SGLD by introducing an additional Metropolis-adjusted step, which provides a faster convergence rate, as shown in Dwivedi et al. (2018); Mangoubi and Vishnoi (2019); Holzmüller and Bach (2023). A more detailed discussion of our implementations of SGLD and MALA can be found in Appendix C.

### 3 GIBBS-BASED INFORMATION CRITERIA

We now develop information-theoretic analyses for Gibbs-based AIC and BIC, following the same classical principles. AIC captures the population risk, and BIC approximates the log marginal likelihood. We demonstrate that our proposed Gibbs-based information criteria align with the classical information criteria in the traditional large  $n$  regime and discuss how our information-theoretic analysis can be generalized to the over-parameterized regime.

#### 3.1 Gibbs-based AIC

As we discussed in Section 2.1, the penalty term in AIC can be viewed as the generalization error of MLE with log-loss. Thus, we start with the following result from Aminian et al. (2021), which provides an exact characterization for the generalization error of the Gibbs algorithm using information measure.

**Proposition 1.** (Aminian et al., 2021) *For the Gibbs algorithm defined in (10), the expected generalization error is*

$$\overline{gen}(P_{\hat{W}|S}^*, P_S) = I_{\text{SKL}}(P_{\hat{W}|S}^*, P_S)/\beta, \quad (13)$$

where  $I_{\text{SKL}}(P_{\hat{W}|S}^*, P_S)$  is the symmetrized KL information between  $\hat{W}$  and  $S$ , defined as follows

$$I_{\text{SKL}}(P_{Y|X}, P_X) \triangleq D(P_{X,Y} \| P_X \otimes P_Y) + D(P_X \otimes P_Y \| P_{X,Y}).$$

Notably, information risk minimization in (8) regularizes the mutual information  $I(\hat{W}; S)$  as a proxy of the generalization error, but the exact generalization error of the Gibbs algorithm is the symmetrized KL information, which is always larger than the mutual information.

As discussed in Section 2.3, we can obtain samples from the Gibbs distribution with SGLD or MALA. Thus, the population risk of the Gibbs algorithm can be

approximated by  $L_P(\hat{W}_{\text{Gibbs}}, P_Z) \approx L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \frac{1}{\beta} I_{\text{SKL}}(P_{\hat{W}|S}^*, P_S)$ , which motivates the following *Gibbs-based AIC*:

$$\text{AIC}^+ \triangleq L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \frac{1}{\beta} I_{\text{SKL}}(P_{\hat{W}|S}^*, P_S). \quad (14)$$

Observe that the penalty term in Gibbs-based AIC is an information measure that captures its generalization error. By investigating the asymptotic behavior of  $I_{\text{SKL}}(P_{\hat{W}|S}^*, P_S)$ , we have the following theorem characterizes the Gibbs-based AIC in the classical asymptotic regime.

**Theorem 1.** (proved in Appendix D) *Consider the log-loss  $\ell(w, \mathbf{z}) = -\log P(y|\mathbf{x}; w)$ , and set  $\beta = n$ . Under proper regularization assumptions in Appendix D, the Gibbs-based AIC has the following form in the regime where  $p$  is fixed and  $n \rightarrow \infty$ :*

$$\text{AIC}^+ = L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \frac{p}{n}. \quad (15)$$

Evidently, our information-theoretic analysis has the same AIC penalty term for the  $\text{AIC}^+$  in the classical regime, which suggests that the generalization error of the Gibbs algorithm (SGLD or MALA) has the same order of  $p/n$  as that of the MLE (SGD) in this regime.

#### 3.2 Gibbs-based BICs

The Gibbs-based BIC is constructed by computing the marginal likelihood  $m(\mathbf{z}^n)$  using the information risk minimization framework. As such, it differs from the standard approach in classical (MLE-based) BIC, as no Laplace approximation is needed.

We now show that the minimum value achieved by the Gibbs algorithm with log-loss in the information risk minimization is the negative log-marginal likelihood.

**Proposition 2.** (proved in Appendix E) *For the Gibbs algorithm  $P_{\hat{W}|S}^*$  defined in (10), if we adopt the log-loss function  $\ell(w, \mathbf{z}) = -\log P(y|\mathbf{x}; w)$ , and set  $\beta = n$ , the marginal likelihood is*

$$\begin{aligned} & -\frac{1}{n} \log m(\mathbf{z}^n) \\ &= \mathbb{E}_{P_{\hat{W}|S=z^n}^*} [L_E(\hat{W}, \mathbf{z}^n)] + \frac{1}{n} D(P_{\hat{W}|S=z^n}^* \| \pi) \quad (16) \\ &= \mathbb{E}_{\pi} [L_E(\hat{W}, \mathbf{z}^n)] - \frac{1}{n} D(\pi \| P_{\hat{W}|S=z^n}^*). \end{aligned}$$

Motivated by Proposition 2, we propose two different Gibbs-based BICs to approximate the marginal likelihood:

$$\text{BIC}^+ \triangleq L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \frac{1}{n} D(P_{\hat{W}|S=z^n}^* \| \pi), \quad (17)$$

$$\text{BIC}^- \triangleq \mathbb{E}_\pi [L_E(W, \mathbf{z}^n)] - \frac{1}{n} D(\pi \| P_{\hat{W}|S=\mathbf{z}^n}^*). \quad (18)$$

Interestingly, given that the empirical risk term is evaluated under the predetermined prior distribution, it is unnecessary to sample from the Gibbs posterior to evaluate  $\text{BIC}^-$ . The  $\text{BIC}^-$  can be directly obtained by computing the second KL divergence term depending on  $P_{\hat{W}|S=\mathbf{z}^n}^*$ .

To compare  $\text{BIC}^+$  with the classical BIC, it suffices to investigate the asymptotic behavior of the KL divergence between the Gibbs posterior distribution and the prior as  $n \rightarrow \infty$ . The final expression of the Gibbs-based  $\text{BIC}^+$  is formally stated in the following Theorem.

**Theorem 2.** (proved in Appendix F) *Under proper regularization assumptions in Appendix F, the Gibbs-based  $\text{BIC}^+$  has the following form in the classical regime where  $p$  is fixed and  $n \rightarrow \infty$ ,*

$$\text{BIC}^+ \triangleq L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \frac{p}{2n} \log n. \quad (19)$$

As expected, we have the same BIC penalty term for the  $\text{BIC}^+$  in the classical regime. The experimental results in Appendix I show that the proposed Gibbs-based AIC and BIC are comparable to their classic counterparts. In the over-parameterized regimes, they are not, as we now develop.

## 4 BICS FOR OVER-PARAMETERIZED RF MODEL

As shown in Figure 1 (right), the classical AIC and BIC fail to capture the double-descent trend exhibited by the test MSE of the RF model. This is due to the fact that the generalization error and marginal likelihood have different behaviors in the over-parameterized regime, which cannot be characterized by the classical asymptotic analysis. The classical analyses heavily rely on the asymptotic normality of MLE and Laplace approximation under certain regularization assumptions, which ignores the prior distribution as  $n \rightarrow \infty$ . Unfortunately, none of these properties hold in the over-parameterized regime where  $p \gg n$ , as there exist an infinite number of possible model parameters that could interpolate  $n$  samples perfectly, making the classical AIC and BIC ill-defined.

However, the Gibbs-based  $\text{AIC}^+$  in (14),  $\text{BIC}^+$  in (17), and  $\text{BIC}^-$  in (18) defined using different information measures can be generalized to over-parameterized models, as Proposition 1 and 2 hold regardless of the values of  $p$  and  $n$ . Since  $\text{AIC}^+$  mainly captures the generalization error of the Gibbs algorithm, which can be

estimated using validation data in practice, our focus lies in extending the analysis of  $\text{BIC}^+$  and  $\text{BIC}^-$  to approximate the marginal likelihood in the over-parameterized regime.

Owing to the complex nature of fitting in the over-parameterized regime, we do not pursue a general asymptotic formula that applies to all model architectures, as in the classic regime. Instead, we refine the Gibbs-based BIC analysis to this regime for the random feature (RF) model.

### 4.1 Random Feature Model

The RF model (Rahimi and Recht, 2008) takes the form of a two-layer neural network with fixed random weights in the first layer. The output of RF model with input data  $\mathbf{x} \in \mathbb{R}^d$  is

$$g(\mathbf{x}) \triangleq \sum_{j=1}^p f\left(\frac{\langle \mathbf{x}, \mathbf{F}_j \rangle}{\sqrt{d}}\right) \mathbf{w}_j = f\left(\frac{\mathbf{x}^\top \mathbf{F}}{\sqrt{d}}\right) \mathbf{w}, \quad (20)$$

where  $\mathbf{w} \in \mathbb{R}^p$  denotes the weights of the model. Moreover,  $\mathbf{F}_j \in \mathbb{R}^d$  denotes the  $j$ th random feature vector, which is the  $j$ th column of the random feature matrix  $\mathbf{F} \in \mathbb{R}^{d \times p}$  whose entries are drawn i.i.d. from  $\mathcal{N}(0, 1)$ . Finally,  $f(\cdot)$  is a point-wise activation function. In our setting, there are  $n$  training samples  $\mathbf{z}^n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , and the  $\mathbf{x}_i$  are drawn i.i.d. from  $\mathcal{N}(0, \mathbf{I}_d)$ .

The parametric distribution family induced by the random feature model takes the form

$$P(y^n | \mathbf{x}^n; \mathbf{w}) \quad (21)$$

$$= \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - f(\frac{\mathbf{x}_i^\top \mathbf{F}}{\sqrt{d}}) \mathbf{w})^2\right),$$

with a fixed random feature matrix  $\mathbf{F}$ , and a Gaussian prior distribution  $\mathbf{w} \sim \mathcal{N}(0, \frac{\sigma^2}{\lambda n} \mathbf{I}_p)$ . The weights of the RF model  $\mathbf{w}$  can be trained using the regularized log-loss

$$\mathcal{L}(\mathbf{w}) = \frac{1}{2\sigma^2} \|\mathbf{Y} - \mathbf{B}\mathbf{w}\|_2^2 + \frac{n}{2} \log(2\pi\sigma^2) + \frac{\lambda n \|\mathbf{w}\|_2^2}{2\sigma^2},$$

where  $\mathbf{B} \triangleq f(\mathbf{X}\mathbf{F}/\sqrt{d}) \in \mathbb{R}^{n \times p}$ , (22)

and we collect the training data in a matrix  $\mathbf{X} \in \mathbb{R}^{n \times d}$  and a vector  $\mathbf{Y} \in \mathbb{R}^n$  to simplify the notation.

As discussed in Mei and Montanari (2022); d'Ascoli et al. (2020), a significant benefit of using the random feature model is that, unlike in standard linear regression, the dimensionality of the input data  $d$  is not entangled with the number of parameters  $p$ .

## 4.2 Gibbs-based BICs for Over-parameterized RF Model

To generalize Gibbs-based BICs to the over-parameterized RF model, it suffices to focus on the second KL-divergence term in (17) and (18). In the random feature model, it can be shown (see Appendix G for details) that the Gibbs algorithm reduces to the Gaussian posterior distribution  $P_{\hat{W}|S=z^n}^* \sim \mathcal{N}(\hat{W}_\lambda, \Sigma_w)$ , with the mean  $\hat{W}_\lambda = (\lambda n \mathbf{I}_p + \mathbf{B}^\top \hat{\mathbf{B}}) \mathbf{1} \mathbf{B}^\top \mathbf{Y}$ , and covariance matrix  $\Sigma_w = \sigma^2 (\lambda n \mathbf{I}_p + \mathbf{B}^\top \mathbf{B})^{-1}$ .

Thus, the KL-divergence between the Gibbs posterior distribution and prior  $\mathcal{N}(0, \frac{\sigma^2}{\lambda n} \mathbf{I}_p)$  is given by

$$D(P_{\hat{W}|S=z^n}^* \|\pi) = \frac{1}{2} \left[ \frac{\lambda n}{\sigma^2} \|\hat{W}_\lambda\|_2^2 + \log \frac{\det(\frac{\sigma^2}{\lambda n} \mathbf{I}_p)}{\det(\Sigma_w)} + \text{tr}(\frac{\lambda n}{\sigma^2} \Sigma_w) - p \right]. \quad (23)$$

And the the KL-divergence between the prior  $\mathcal{N}(0, \frac{\sigma^2}{\lambda n} \mathbf{I}_p)$  and Gibbs distribution is given by

$$D(\pi \| P_{\hat{W}|S=z^n}^*) = \frac{1}{2} \left[ \hat{W}_\lambda^\top (\Sigma_w)^{-1} \hat{W}_\lambda - \log \frac{\det(\frac{\sigma^2}{\lambda n} \mathbf{I}_p)}{\det(\Sigma_w)} + \text{tr}(\frac{\sigma^2}{\lambda n} \Sigma_w^{-1}) - p \right]. \quad (24)$$

To obtain a convenient expression for the remaining determinant and trace terms, we first impose restrictions on the activation function  $f(\cdot)$ . Therefore, these two terms can be characterized using random matrix theory by studying the eigenvalues of  $\Sigma \triangleq \mathbf{B}^\top \mathbf{B} / (\lambda n) + \mathbf{I}_p$  in the over-parameterized regime. In particular, for activation functions  $f(\cdot)$  that satisfy conditions

$$\begin{aligned} \mathbb{E}[f(\varepsilon)] = 0, \quad \mathbb{E}[f(\varepsilon)^2] = 1, \quad \mathbb{E}[f'(\varepsilon)] = 0, \\ |\mathbb{E}[f(\varepsilon)^k]| < \infty, \quad \text{for } k > 1, \end{aligned} \quad (25)$$

where  $\varepsilon \sim \mathcal{N}(0, 1)$ , the following theorem characterizes the KL divergence term in the over-parameterized RF model.

**Theorem 3.** *For activation functions  $f(\cdot)$  satisfying the conditions in (25), as  $n, d, p \rightarrow \infty$  with  $p/d \rightarrow r_1$ ,  $n/d \rightarrow r_2$ , and  $r_1/r_2 = r$ , where  $r_1, r_2 \in (0, \infty)$ , we have*

$$\begin{aligned} \frac{1}{n} D(P_{\hat{W}|S=z^n}^* \|\pi) \\ \rightarrow \frac{\lambda}{2\sigma^2} \|\hat{W}_\lambda\|_2^2 - \frac{\lambda}{8} \mathcal{F}\left(\frac{1}{\lambda}, r\right) + \frac{1}{2} V(1/\lambda, r) \end{aligned} \quad (26)$$

almost surely, where

$$\begin{aligned} V(\gamma, r) \triangleq r \log \left( 1 + \gamma - \frac{1}{4} \mathcal{F}(\gamma, r) \right) - \frac{1}{4\gamma} \mathcal{F}(\gamma, r) \\ + \log \left( 1 + \gamma r - \frac{1}{4} \mathcal{F}(\gamma, r) \right), \end{aligned} \quad (27)$$

$$\mathcal{F}(\gamma, r) \triangleq \left( \sqrt{\gamma(1+\sqrt{r})^2 + 1} - \sqrt{\gamma(1-\sqrt{r})^2 + 1} \right)^2.$$

*Sketch of Proof.* The proof is based on the results from Pennington and Worah (2017), which shows that the distribution of the eigenvalues of the random matrix  $\mathbf{B}^\top \mathbf{B} / n$  converges to the Marchenko-Pastur distribution with shape parameter  $r$  (a well-studied distribution in random matrix theory (Tulino et al., 2004)). The detailed proof is provided in Appendix H.  $\square$

**Remark 1.** *An example of an activation function that satisfies all the assumptions we made in (25) is  $f(x) = (x^2 - 1)/\sqrt{2}$ . More examples of such activation functions can be found in Pennington and Worah (2017). We further note that the assumption  $\mathbb{E}[f'(\varepsilon)] = 0$  on the activation function is used only to obtain a simple closed-form for the KL divergence, a more general result by considering the Stieltjes transform of  $\mathbf{B}^\top \mathbf{B} / n$  for other activation functions is provided in Appendix H.*

Theorem 3 motivates us to define the following Gibbs-based BICs for the over-parameterized RF model to approximate the marginal likelihood,

$$\begin{aligned} \text{BIC}^+ \triangleq L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \underbrace{\frac{\lambda}{2\sigma^2} \|\hat{W}_\lambda\|_2^2}_{\ell_2 \text{ term}} \\ - \underbrace{\frac{\lambda}{8} \mathcal{F}\left(\frac{1}{\lambda}, r\right) + \frac{1}{2} V(1/\lambda, r)}_{\text{covariance term}} \\ \text{BIC}^- \triangleq \mathbb{E}_\pi [L_E(W, \mathbf{z}^n)] - \frac{1}{2n} \hat{W}_\lambda^\top (\Sigma_w)^{-1} \hat{W}_\lambda \\ - \frac{1}{2n} \text{tr}(\mathbf{I}_p + \frac{1}{\lambda n} \mathbf{B}^\top \mathbf{B}) + \frac{1}{2} V(1/\lambda, r) + \frac{p}{2n}. \end{aligned} \quad (28)$$

Depending on the specific sampling method employed, as detailed in Section 2.3, the term  $L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n)$  can be estimated with either  $L_E(\hat{W}_{\text{SGLD}}, \mathbf{z}^n)$  or  $L_E(\hat{W}_{\text{MALA}}, \mathbf{z}^n)$ . When  $\sigma$  is comparatively small and  $\lambda$  is comparatively large,  $\|\hat{W}_\lambda\|_2^2$  can be substituted by  $\|\hat{W}_{\text{SGLD}}\|_2^2$  or  $\|\hat{W}_{\text{MALA}}\|_2^2$ , as the posterior is primarily driven by the mean. The penalty term of  $\text{BIC}^+$  consists of the  $\ell_2$  norm of the learned weights, and two other terms capture the log determinant and trace in the over-parameterized regime, which will be referred to as the covariance term altogether in the next section.

## 5 EXPERIMENT AND DISCUSSION

In this section, we instantiate a two-layer RF model described in (20), where the first layer is designated for feature mapping and is kept random, and we only train the parameter in the second layer. See Appendix I for experimental details and additional results.

We evaluate the over-parameterized Gibbs-based  $\text{BIC}^+$  and  $\text{BIC}^-$  using  $n = 200$  samples generated by the

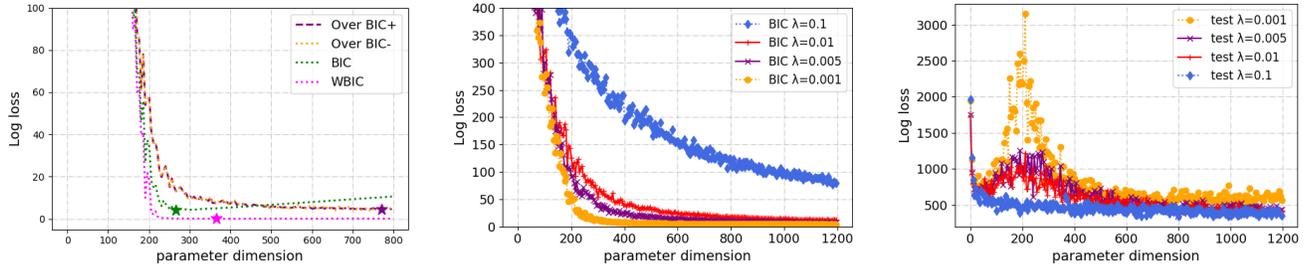


Figure 2: A comparison between different BICs in over-parameterized RF model when  $\lambda = 0.001$  (left); A comparison between  $\text{BIC}^+$  (middle) and population risk (right) with varying  $\lambda$ .

linear model

$$y_i = \mathbf{x}_i^\top \mathbf{w}^* + \epsilon_i, \quad \mathbf{w}^* \in \mathbb{R}^d, \|\mathbf{w}^*\|_2^2 = 1, \epsilon_i \sim \mathcal{N}(0, \sigma_d^2),$$

with input dimension  $d = 400$ , noise  $\sigma_d^2 = 0.1$ , and we use MALA to sample from Gibbs distribution<sup>1</sup>.

**Selection of  $\sigma^2$ .** By examining the mean of the Gibbs posterior, given as  $\hat{W}_\lambda = (\lambda n \mathbf{I}_p + \mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{Y}$ , and its covariance  $\Sigma_w = \sigma^2 (\lambda n \mathbf{I}_p + \mathbf{B}^\top \mathbf{B})^{-1}$ , it becomes evident that  $\sigma^2$  impacts only the variance of the posterior and not its mean. A large  $\sigma^2$  can introduce instability in the sampling from the posterior and deteriorate the test performance. Thus, in our configuration, we opt for a smaller  $\sigma = 0.05$  to ensure good model performance, even if it results in a larger scaling factor for empirical risk. In our upcoming discussion, We will investigate how the other parameter  $\lambda$  in the prior distribution influences the posterior.

**Double-descent of population risk.** As depicted in Figure 1 (left), the peak of test loss is located at the interpolation threshold, i.e., when  $p = n = 200$ , resulting in the highest generalization error. As  $p$  continues to increase, the test error begins to decline again, even falling below the levels observed in the under-parameterized regime  $p < n$ .

**Comparison of different BICs.** In the right panel of Figure 1 and Figure 2 (left), it is evident that the classic BIC prefers moderate model due to the ill-defined issue we discussed before. In Figure 2 (left), we also compare our Gibbs-based BICs with the WBIC proposed in Watanabe (2013). The analysis of WBIC is based on approximating the marginal likelihood for singular models in the classical large  $n$  regime, and such an approximation becomes inaccurate in the considered over-parameterized regime. Thus, WBIC also fails to capture the marginal likelihood, whereas our over-parameterized  $\text{BIC}^+$  and  $\text{BIC}^-$  in (28) and (29)

based on the exact forms of the marginal likelihood, succeeds in selecting the large model.

**Mismatch between BIC and population risk.** Note that even though the Gibbs-based BICs might favor a larger model for increased values of  $p$ , it does not exhibit double-descent behavior. This discrepancy becomes evident when examining Figure 2 (middle) and Figure 2 (right), where BIC exhibits a distinctive pattern compared to the population risk. A similar mismatch between marginal likelihood and generalization for ERM has been observed in Lotfi et al. (2022).

We further investigate the inconsistency between the marginal likelihood and population risk for the Gibbs algorithm in Figure 3. Unlike the classical BIC, where the penalty term  $p \log n / (2n)$  is order-wise larger than the  $p/n$  term in classical AIC, it can be seen that the KL divergence term in the over-parameterized  $\text{BIC}^+$  can be significantly smaller than the generalization error  $I_{\text{SKL}}$  in Figure 3, depending on the value of  $\lambda$  and  $p$ . Thus, the mismatch between marginal likelihood (BIC) and population risk (AIC) is even more complicated in the over-parameterized setting due to the influence of prior distribution.

**Similarity between KL divergence and generalization.** When comparing the KL divergence term of  $\text{BIC}^+$  in Figure 3 (left) and the generalization error in Figure 3 (right) with the same  $\lambda$ , a similar trend emerges as  $p$  increases. To understand such similarity between the two terms, we decompose the penalty term of the over-parameterized  $\text{BIC}^+$  in Figure 4 into  $\ell_2$  term, covariance term.

As shown in Figure 4, when  $p \leq n$ , the model prior can be ignored, and the training loss becomes the dominant factor of  $\text{BIC}^+$  (training loss is plotted in Appendix I due its large scale). In this case, the KL divergence and the covariance term increase with  $p$ , corresponding to the classical BIC. When  $p \geq n$ , the KL divergence penalty term dominates the over-parameterized  $\text{BIC}^+$ . In this regime, multiple weights can fit the training

<sup>1</sup>The code of our experiment is available at <https://github.com/HaoboChen11/Gibbs-Based-Information-Criteria-and-the-Over-Parameterized-Regime>.

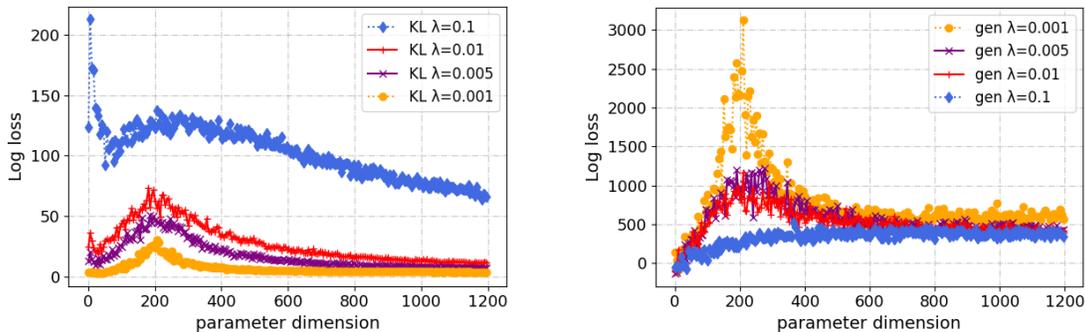


Figure 3: A comparison between the KL-divergence term in  $BIC^+$  (left) and the generalization error term in  $AIC^+$  (right) with varying  $\lambda$ . Note the different scales in the vertical axis.

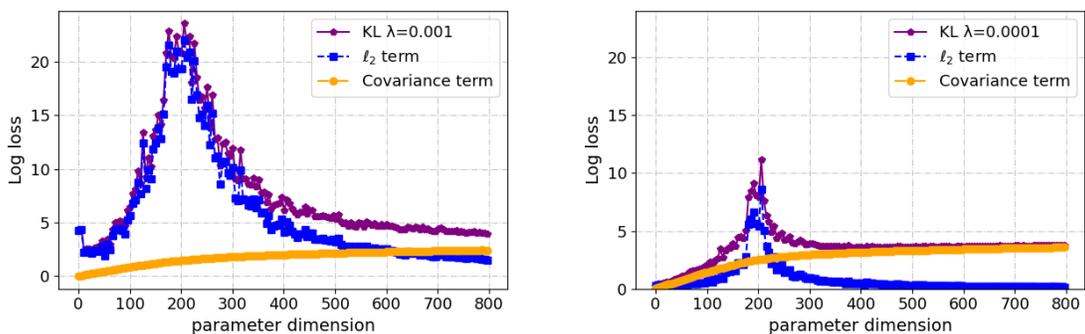


Figure 4: A decomposition of the terms in over-parameterized  $BIC^+$  in (28) with  $\lambda = 0.001$  (left), and  $\lambda = 0.0001$  (right).

data perfectly. From Figure 4, we note that regardless of the chosen  $\lambda$ , the  $\ell_2$  term exhibits double-descent behavior and decreases as  $p$  increases. This suggests that the Gibbs algorithm prefers the weights with smaller  $\ell_2$  norm to fit the data. Note that similar phenomena are observed for SGD, and generalization error bounds using different weights norms are established in Neyshabur et al. (2017); Bartlett et al. (2017), i.e., a smaller  $\ell_2$  norm can lead to a better generalization performance. Thus, the behavior of the  $\ell_2$  norm may elucidate why, for a given  $\lambda$  value, both KL divergence and the generalization error  $I_{SKL}$  exhibit similar double-descent behavior.

**Mismatch between KL divergence and generalization.** The comparison between Figure 3 (left) and Figure 3 (right) with different values of  $\lambda$  reveals a mismatch between KL divergence and generalization error. A closer examination of Figure 4 (left) and Figure 4 (right) shows that larger  $\lambda$  results in a larger  $\ell_2$  term in KL divergence. Note that  $\|\hat{W}_\lambda\|_2^2$  is increasing with smaller  $\lambda$ , but the  $\ell_2$  term in KL divergence equals to  $\frac{\lambda}{2\sigma^2}\|\hat{W}_\lambda\|_2^2$ . Consequently, the KL divergence places a more substantial penalty on larger values of  $\lambda$ , which explains the mismatch between the KL divergence and  $I_{SKL}$  with different priors by varying  $\lambda$ . A comprehen-

sive discussion can be found in the experimental results section, Appendix I.

## 6 CONCLUDING REMARKS

In this paper, we introduce novel information criteria based on the Gibbs algorithm to capture the double-descent behavior observed in over-parameterized models, thereby extending the analysis of classical AIC and BIC to modern regimes. Demonstrated to be effective in selecting suitable models within over-parameterized regimes, our criteria reveal that the mismatch between the log marginal likelihood and the test error can be captured by the disparity between two information measures. Interestingly, such a phenomenon highly depends on the selection of the prior distributions, which is quite different from traditional understanding in the classical regime.

Importantly, our Gibbs-based AIC and BICs, grounded in information measures, can adapt to a broad spectrum of models, beyond asymptotic regimes. Therefore, a promising future direction is to extend our analysis to other models, e.g., multi-layer RF models, which bear a closer resemblance to deep neural networks.

## References

- Advani, M. S., Saxe, A. M., and Sompolinsky, H. (2020). High-dimensional dynamics of generalization error in neural networks. *Neural Networks*.
- Akaike, H. (1981). Likelihood of a model and information criteria. *Journal of Econometrics*, 16(1):3–14.
- Aminian, G., Bu, Y., Toni, L., Rodrigues, M., and Wornell, G. (2021). An exact characterization of the generalization error for the Gibbs algorithm. *Advances in Neural Information Processing Systems*, 34:8106–8118.
- Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. (2017). Spectrally-normalized margin bounds for neural networks. *Advances in neural information processing systems*, 30.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. (2020). Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*.
- Belkin, M., Hsu, D., Ma, S., and Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854.
- Belkin, M., Hsu, D., and Xu, J. (2020). Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180.
- Chen, J. and Chen, Z. (2008). Extended Bayesian information criteria for model selection with large model spaces. *Biometrika*, 95(3):759–771.
- d’Ascoli, S., Sagun, L., and Biroli, G. (2020). Triple descent and the two kinds of overfitting: Where & why do they appear? *Advances in Neural Information Processing Systems*, 33:3058–3069.
- Deng, Z., Kammoun, A., and Thrampoulidis, C. (2022). A model of double descent for high-dimensional binary linear classification. *Information and Inference: A Journal of the IMA*, 11(2):435–495.
- Dwivedi, R., Chen, Y., Wainwright, M. J., and Yu, B. (2018). Log-concave sampling: Metropolis-hastings algorithms are fast! In *Conference on learning theory*, pages 793–797. PMLR.
- Dwivedi, R., Singh, C., Yu, B., and Wainwright, M. J. (2020). Revisiting complexity and the bias-variance tradeoff. *arXiv preprint arXiv:2006.10189*.
- Dziugaite, G. K. and Roy, D. (2018). Entropy-SGD optimizes the prior of a PAC-Bayes bound: Generalization properties of Entropy-SGD and data-dependent priors. In *International Conference on Machine Learning*, pages 1377–1386. PMLR.
- d’Ascoli, S., Refinetti, M., Biroli, G., and Krzakala, F. (2020). Double trouble in double descent: Bias and variance (s) in the lazy regime. In *International Conference on Machine Learning*, pages 2280–2290. PMLR.
- Fan, Y. and Tang, C. Y. (2013). Tuning parameter selection in high dimensional penalized likelihood. *Journal of the Royal Statistical Society: SERIES B: Statistical Methodology*, pages 531–552.
- Geiger, M., Spigler, S., d’Ascoli, S., Sagun, L., Baity-Jesi, M., Biroli, G., and Wyart, M. (2019). Jamming transition as a paradigm to understand the loss landscape of deep neural networks. *Physical Review E*, 100(1):012115.
- Geman, S., Bienenstock, E., and Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1):1–58.
- Gerace, F., Loureiro, B., Krzakala, F., Mézard, M., and Zdeborová, L. (2020). Generalisation error in learning with random features and the hidden manifold model. In *International Conference on Machine Learning*, pages 3452–3462. PMLR.
- Germain, P., Bach, F., Lacoste, A., and Lacoste-Julien, S. (2016). PAC-Bayesian theory meets Bayesian inference. *Advances in Neural Information Processing Systems*, 29.
- Gibbs, J. W. (1902). Elementary principles of statistical mechanics. *Compare*, 289:314.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. (2022). Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949–986.
- Heckel, R. and Yilmaz, F. F. (2020). Early stopping in deep networks: Double descent and how to eliminate it. *arXiv preprint arXiv:2007.10099*.
- Hodgkinson, L., van der Heide, C., Salomone, R., Roosta, F., and Mahoney, M. W. (2023). The interpolating information criterion for overparameterized models. *arXiv preprint arXiv:2307.07785*.
- Holzmüller, D. and Bach, F. (2023). Convergence rates for non-log-concave sampling and log-partition estimation. *arXiv preprint arXiv:2303.03237*.
- Immer, A., Bauer, M., Fortuin, V., Rätsch, G., and Emtiyaz, K. M. (2021). Scalable marginal likelihood estimation for model selection in deep learning. In *International Conference on Machine Learning*, pages 4563–4573. PMLR.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical review*, 106(4):620.
- Kini, G. R. and Thrampoulidis, C. (2020). Analytic study of double descent in binary classification: The impact of loss. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2527–2532. IEEE.

- Kleijn, B. J., van der Vaart, A. W., et al. (2012). The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381.
- Liu, F., Suykens, J., and Cevher, V. (2022). On the double descent of random features models trained with SGD. *Advances in Neural Information Processing Systems*, 35:34966–34980.
- Lotfi, S., Izmailov, P., Benton, G., Goldblum, M., and Wilson, A. G. (2022). Bayesian model selection, the marginal likelihood, and generalization. In *International Conference on Machine Learning*, pages 14223–14247. PMLR.
- Mangoubi, O. and Vishnoi, N. K. (2019). Nonconvex sampling with the Metropolis-adjusted Langevin algorithm. In *Conference on Learning Theory*, pages 2259–2293. PMLR.
- Mei, S. and Montanari, A. (2022). The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766.
- Murphy, K. P. (2007). Conjugate Bayesian analysis of the Gaussian distribution. *def*, 1(2 $\sigma$ 2):16.
- Muthukumar, V., Vodrahalli, K., Subramanian, V., and Sahai, A. (2020). Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*.
- Nakkiran, P. (2019). More data can hurt for linear regression: Sample-wise double descent. *arXiv preprint arXiv:1912.07242*.
- Nakkiran, P., Kaplun, G., Bansal, Y., Yang, T., Barak, B., and Sutskever, I. (2019). Deep double descent: Where bigger models and more data hurt. In *International Conference on Learning Representations*.
- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. (2017). Exploring generalization in deep learning. *Advances in neural information processing systems*, 30.
- Palomar, D. P. and Verdú, S. (2008). Lautum information. *IEEE transactions on information theory*, 54(3):964–975.
- Pennington, J. and Worah, P. (2017). Nonlinear random matrix theory for deep learning. In *Advances in Neural Information Processing Systems*, pages 2637–2646.
- Perlaza, S. M., Bisson, G., Esnaola, I., Jean-Marie, A., and Rini, S. (2022). Empirical risk minimization with relative entropy regularization: Optimality and sensitivity analysis. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pages 684–689. IEEE.
- Perlaza, S. M., Esnaola, I., Bisson, G., and Poor, H. V. (2023). On the validation of Gibbs algorithms: Training datasets, test datasets and their aggregation. *arXiv preprint arXiv:2306.12380*.
- Raginsky, M., Rakhlin, A., and Telgarsky, M. (2017). Non-convex learning via stochastic gradient Langevin dynamics: a nonasymptotic analysis. In *Conference on Learning Theory*, pages 1674–1703. PMLR.
- Rahimi, A. and Recht, B. (2008). Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184.
- Schwarz, G. et al. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464.
- Tulino, A. M., Verdú, S., and Verdu, S. (2004). *Random matrix theory and wireless communications*. Now Publishers Inc.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- Watanabe, S. (2013). A widely applicable Bayesian information criterion. *The Journal of Machine Learning Research*, 14(1):867–897.
- Weinan, E., Ma, C., and Wu, L. (2020). A comparative analysis of optimization and generalization properties of two-layer neural network and random feature models under gradient descent dynamics. *Science China Mathematics*, pages 1–24.
- Xu, A. and Raginsky, M. (2017). Information-theoretic analysis of generalization capability of learning algorithms. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pages 2524–2533.
- Yang, Z., Yu, Y., You, C., Steinhardt, J., and Ma, Y. (2020). Rethinking bias-variance trade-off for generalization of neural networks. In *International Conference on Machine Learning*, pages 10767–10777. PMLR.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. In *Proc. International Conference on Learning Representations (ICLR)*.
- Zhang, T. (2006). Information-theoretic upper and lower bounds for statistical estimation. *IEEE Transactions on Information Theory*, 52(4):1307–1321.

## Checklist

1. For all models and algorithms presented, check if you include:
  - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes] (See Sections 2 and 4)
  - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes] (See Appendix I)
  - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [No] (We will submit the source code during the camera-ready phase.)
2. For any theoretical claim, check if you include:
  - (a) Statements of the full set of assumptions of all theoretical results. [Yes] (See Sections 3 and 4)
  - (b) Complete proofs of all theoretical results. [Yes] (All proofs are in the Appendix with appropriate references in Sections 3 and 4.)
  - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
  - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes] (Please check Section 5 and Appendix I.)
  - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes] (Please check Appendix I.)
  - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes] (Please check Appendix I.)
  - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes] (Please check Appendix I.)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
  - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
  - (b) The license information of the assets, if applicable. [Not Applicable]
  - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
  - (d) Information about consent from data providers/curators. [Not Applicable]
  - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
  - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
  - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
  - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

## Supplementary Materials

---

### A DERIVATION OF CLASSICAL AIC

We start by formally presenting the regularization conditions required for the standard asymptotic normality of MLE in the classical regime.

**Assumption 1. Regularity Conditions for MLE**

1. *Identifiability:*  $P(y|\mathbf{x};\boldsymbol{\theta}) \neq P(y|\mathbf{x};\boldsymbol{\theta}')$  for  $\boldsymbol{\theta} \neq \boldsymbol{\theta}'$ .
2.  $\Theta$  is an open subset of  $\mathbb{R}^p$ .
3. The function  $\log P(y|\mathbf{x};\boldsymbol{\theta})$  is three times continuously differentiable with respect to  $\boldsymbol{\theta}$ .
4. There exist functions  $F_1(z) : \mathcal{Z} \rightarrow \mathbb{R}, F_2(z) : \mathcal{Z} \rightarrow \mathbb{R}$  and  $M(z) : \mathcal{Z} \rightarrow \mathbb{R}$ , such that

$$\mathbb{E}_{Z \sim P(\mathbf{z};\boldsymbol{\theta})} [M(Z)] < \infty,$$

and the following inequalities hold for any  $\boldsymbol{\theta} \in \Theta$ ,

$$\begin{aligned} \left| \frac{\partial \log P(y|\mathbf{x};\boldsymbol{\theta})}{\partial \theta_i} \right| &< F_1(z), & \left| \frac{\partial^2 \log P(y|\mathbf{x};\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right| &< F_1(z), \\ \left| \frac{\partial^3 \log P(y|\mathbf{x};\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j \partial \theta_k} \right| &< M(z), & i, j, k = 1, 2, \dots, p. \end{aligned}$$

5. The following inequality holds for an arbitrary  $\boldsymbol{\theta} \in \Theta$  and  $i, j = 1, 2, \dots, p$ ,

$$0 < \mathbb{E} \left[ \frac{\partial \log P(y|\mathbf{x};\boldsymbol{\theta})}{\partial \theta_i} \frac{\partial \log P(y|\mathbf{x};\boldsymbol{\theta})}{\partial \theta_j} \right] < \infty.$$

In the following, we provide proof of the classical AIC for reference.

The AIC model selection in (4) is equivalent to:

$$\operatorname{argmin}_k \mathbb{E}_{P_Z} [-\log P_k(y|\mathbf{x};\hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)})] = \operatorname{argmin}_k \mathbb{E}_{P_S} [L_E(\hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)}, S)] + \overline{ge\bar{n}}(\hat{\boldsymbol{\theta}}_{\text{ML}}^{(k)}, P_Z). \quad (30)$$

As  $n \rightarrow \infty$ , under the above regularization conditions, which guarantee that  $\hat{\boldsymbol{\theta}}_{\text{ML}}$  is unique, the asymptotic normality of the MLE states that the distribution of  $\hat{\boldsymbol{\theta}}_{\text{ML}}$  converges to

$$\mathcal{N}(\boldsymbol{\theta}^*, \frac{1}{n} J(\boldsymbol{\theta}^*)^{-1} I(\boldsymbol{\theta}^*) J(\boldsymbol{\theta}^*)^{-1}), \text{ with } \boldsymbol{\theta}^* \triangleq \operatorname{argmin}_{\boldsymbol{\theta} \in \Theta} D(P_Z \| P(y|\mathbf{x};\boldsymbol{\theta})), \quad (31)$$

where

$$J(\boldsymbol{\theta}) \triangleq \mathbb{E}_{P_Z} [-\nabla_{\boldsymbol{\theta}}^2 \log P(y|\mathbf{x};\boldsymbol{\theta})] \text{ and } I(\boldsymbol{\theta}) \triangleq \mathbb{E}_{P_Z} [\nabla_{\boldsymbol{\theta}} \log P(y|\mathbf{x};\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \log P(y|\mathbf{x};\boldsymbol{\theta})^\top]. \quad (32)$$

When the true model is in the parametric family  $P_Z = P(y|\mathbf{x};\boldsymbol{\theta}^*)$ , we have  $J(\boldsymbol{\theta}^*) = I(\boldsymbol{\theta}^*)$ , which is the Fisher information matrix.

Thus, the generalization term can be written as

$$\begin{aligned} -\overline{ge\bar{n}}(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z) &= \mathbb{E}_{P_S} [L_E(\hat{\boldsymbol{\theta}}_{\text{ML}}, S)] - L_P(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z) \\ &= \mathbb{E}_{P_S} [L_E(\hat{\boldsymbol{\theta}}_{\text{ML}}, S)] - \mathbb{E}_{P_S} [L_E(\boldsymbol{\theta}^*, S)] + \mathbb{E}_{P_S} [L_E(\boldsymbol{\theta}^*, S)] - L_P(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z) \\ &= \mathbb{E}_{P_S} [L_E(\hat{\boldsymbol{\theta}}_{\text{ML}}, S) - L_E(\boldsymbol{\theta}^*, S)] + L_P(\boldsymbol{\theta}^*, P_Z) - L_P(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z). \end{aligned} \quad (33)$$

As  $\hat{\boldsymbol{\theta}}_{\text{ML}}$  minimizes  $L_E(\hat{\boldsymbol{\theta}}_{\text{ML}}, S)$ , we take the Taylor expansion of  $L_E(\boldsymbol{\theta}^*, S)$  around the point  $\hat{\boldsymbol{\theta}}_{\text{ML}}$

$$L_E(\boldsymbol{\theta}^*, S) = L_E(\hat{\boldsymbol{\theta}}_{\text{ML}}, S) + \frac{1}{2}(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top J(\hat{\boldsymbol{\theta}}_{\text{ML}})(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{\text{ML}}) + \dots \quad (34)$$

And the Taylor expansion of  $L_P(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z)$  around  $\boldsymbol{\theta}^*$  yields

$$L_P(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z) = L_P(\boldsymbol{\theta}^*, P_Z) + \frac{1}{2}(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top J(\boldsymbol{\theta}^*)(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{\text{ML}}) + \dots \quad (35)$$

If we use the quadratic approximation of (34) and (35) in (33), we can get the following asymptotic expression for the generalization error

$$\overline{gen}(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z) = \frac{1}{n} \text{tr}(I(\boldsymbol{\theta}^*)J(\boldsymbol{\theta}^*)^{-1}), \quad (36)$$

where the last step is due to the asymptotic normality of the MLE, and  $J(\hat{\boldsymbol{\theta}}_{\text{ML}})$  will converge to its expectation by the consistency of MLE and the strong law of large numbers. For the case where the true model is in the parametric family  $P_Z = P(y|\mathbf{x}; \boldsymbol{\theta}^*)$ ,  $J(\boldsymbol{\theta}^*) = I(\boldsymbol{\theta}^*)$ ,  $\overline{gen}(\hat{\boldsymbol{\theta}}_{\text{ML}}, P_Z) = p/n$ . Thus, plug the asymptotic generalization error term back in (30) can get,

$$\text{AIC} = -\frac{\hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}})}{n} + \frac{p}{n}. \quad (37)$$

## B DERIVATION OF CLASSICAL BIC

The Taylor expansion of the log-likelihood function  $\hat{L}(\boldsymbol{\theta})$  around  $\hat{\boldsymbol{\theta}}_{\text{ML}}$  yields

$$\hat{L}(\boldsymbol{\theta}) = \hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}}) - \frac{n}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top J(\hat{\boldsymbol{\theta}}_{\text{ML}})(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}}) + \dots \quad (38)$$

Similarly, we can expand the prior distribution  $\pi(\boldsymbol{\theta})$  with Taylor series around  $\hat{\boldsymbol{\theta}}_{\text{ML}}$  as

$$\pi(\boldsymbol{\theta}) = \pi(\hat{\boldsymbol{\theta}}_{\text{ML}}) + (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top \nabla \pi(\boldsymbol{\theta})|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}_{\text{ML}}} + \dots \quad (39)$$

The Laplace approximation takes advantage of the fact that when  $n$  is sufficiently large, the integrand is concentrated in a neighborhood of the mode of  $\hat{L}(\boldsymbol{\theta})$ , i.e., the maximum likelihood (ML) estimator  $\hat{\boldsymbol{\theta}}_{\text{ML}}$ . Thus,

$$\begin{aligned} m(\mathbf{z}^n) &= \int \exp\{\hat{L}(\boldsymbol{\theta})\} \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &\approx \int \exp\left\{\hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}}) - \frac{n}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top J(\hat{\boldsymbol{\theta}}_{\text{ML}})(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})\right\} \left(\pi(\hat{\boldsymbol{\theta}}_{\text{ML}}) + (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top \nabla \pi(\boldsymbol{\theta})|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}_{\text{ML}}}\right) d\boldsymbol{\theta} \\ &\stackrel{(a)}{\approx} \exp\{\hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}})\} \pi(\hat{\boldsymbol{\theta}}_{\text{ML}}) \int \exp\left\{-\frac{n}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top J(\hat{\boldsymbol{\theta}}_{\text{ML}})(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})\right\} d\boldsymbol{\theta} \\ &= \exp\{\hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}})\} \pi(\hat{\boldsymbol{\theta}}_{\text{ML}}) (2\pi)^{p/2} n^{-p/2} |J(\hat{\boldsymbol{\theta}}_{\text{ML}})|^{-1/2}, \end{aligned}$$

where (a) follows from the following fact

$$\int (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top \exp\left\{-\frac{n}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})^\top J(\hat{\boldsymbol{\theta}}_{\text{ML}})(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{\text{ML}})\right\} d\boldsymbol{\theta} = 0. \quad (40)$$

Taking the logarithm of this expression and multiplying it by  $-\frac{1}{n}$ , we obtain

$$-\frac{1}{n} \log(m(\mathbf{z}^n)) \approx -\frac{1}{n} \hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}}) + \frac{p}{2n} \log \frac{n}{2\pi} + \frac{1}{n} \log |J(\hat{\boldsymbol{\theta}}_{\text{ML}})| - \frac{1}{n} \log \pi(\hat{\boldsymbol{\theta}}_{\text{ML}}). \quad (41)$$

Note that  $J(\hat{\boldsymbol{\theta}}_{\text{ML}})$  is a random Hessian matrix. As  $n \rightarrow \infty$ ,  $J(\hat{\boldsymbol{\theta}}_{\text{ML}})$  will converge to its expectation by the consistency of MLE and the strong law of large numbers. Thus, both  $\log |J(\hat{\boldsymbol{\theta}}_{\text{ML}})|$  and  $\log \pi(\hat{\boldsymbol{\theta}}_{\text{ML}})$  have order less than  $O(1)$  with respect to the sample size  $n$ , and can be ignored in the BIC, which gives the following classical form of BIC:

$$\text{BIC} = -\frac{\hat{L}(\hat{\boldsymbol{\theta}}_{\text{ML}})}{n} + \frac{p \log n}{2n}. \quad (42)$$

## C DETAILS OF SGLD AND GIBBS ALGORITHM

In Raginsky et al. (2017), it is shown that the following SGLD updates

$$\hat{W}_{t+1} = \hat{W}_t - \eta \nabla F_Z(W) + \sqrt{\frac{2\eta}{\beta}} \zeta_t, \quad t = 0, 1, \dots, \quad (43)$$

will converge to the Gibbs algorithm satisfying  $P_{\hat{W}|Z} \propto \exp(-\beta F_Z(\hat{W}))$ .

A detailed description of Metropolis adjusted Langevin algorithm is provided in Algorithm C.

---

### Algorithm 1 Metropolis adjusted Langevin algorithm (MALA)

---

**Require:** Step size  $\eta$  and a sample  $w_0$  from a starting distribution  $\mu_0$

**Ensure:** Sequence  $w_1, w_2, \dots$

```

1: for  $i = 0, 1, \dots$  do
2:    $w_{i+1} \sim N\left(w_i - \eta \nabla F_Z(w_i), \frac{2\eta}{\beta} I_d\right)$  % Propose a new state
3:    $\alpha_{i+1} = \min\left\{1, \frac{\exp(-F_Z(w_{i+1}) - \beta \|w_i - w_{i+1} + \eta \nabla F_Z(w_{i+1})\|_2^2 / 4\eta)}{\exp(-F_Z(w_i) - \beta \|w_{i+1} - w_i + \eta \nabla F_Z(w_i)\|_2^2 / 4\eta)}\right\}$ 
4:    $U_{i+1} \sim U[0, 1]$ 
5:   if  $U_{i+1} \leq \alpha_{i+1}$  then
6:      $w_{i+1} \leftarrow w_{i+1}$  % accept the sample
7:      $w_{i+1} \leftarrow w_i$  % reject the sample
8:   end if
9: end for

```

---

Note that our Gibbs algorithm contains an arbitrary prior distribution  $\pi$ , i.e.,  $P_{\hat{W}|S} \propto \pi(\hat{W}) \exp^{-nL_E(\hat{W}, S)}$ . To have the same format as  $F_Z(W)$ , we let

$$F_Z(w) = L_E(w, s) - \frac{1}{n} \log \pi(w). \quad (44)$$

Thus, the loss function used in the SGLD or MALA update becomes

$$\ell(w, z_i) = -\log P(y_i | x_i, w) - \frac{1}{n} \log \pi(w). \quad (45)$$

When the prior follows a Gaussian distribution  $\pi(w) \sim \mathcal{N}(0, \frac{\sigma^2}{\lambda n} I_p)$ , the second term in the loss function can be viewed as a regularization term derived from the log prior. It is crucial to notice that the empirical log loss in the subsequent SGLD or MALA computation only includes the  $L_E(w, s)$  term and does not incorporate the regularization term. As shown in Raginsky et al. (2017), the regularizer term does not violate the assumptions used in the loss function, ensuring that our loss function will also converge to the desired Gibbs distribution.

In Figure 5 and Figure 6, when examining both the training error and the learned parameter norm over different epochs, we observe that both algorithms converge to a similar value. However, it is noteworthy that MALA exhibits a faster convergence in terms of the number of epochs, as suggested by existing theoretical results.

When the number of training samples  $n$  is smaller, as in Figure 5, despite the computational overhead introduced by the additional Metropolis-adjusted step, MALA still outpaces SGLD by approximately 3 seconds. However, for larger sample sizes in Figure 6, this Metropolis-adjusted step incurs more computational time. Consequently, MALA's convergence time slightly exceeds SGLD's by 0.4 seconds. Thus, in our experiments, we recommend using MALA for smaller sample sizes, while for larger samples, like when  $n = 600$  in the Classic regime, SGLD is more computationally efficient.

## D PROOF OF THEOREM 1

We note that the original idea of this result is coming from the discussion in Section 4 of Aminian et al. (2021).

We further assume that the parametric family  $\{P(z|\mathbf{w}), \mathbf{w} \in \mathcal{W}\}$  and prior  $\pi(\mathbf{w})$  satisfy all the regularization conditions required for the Bernstein–von-Mises theorem (Van der Vaart, 2000) and the asymptotic Normality of

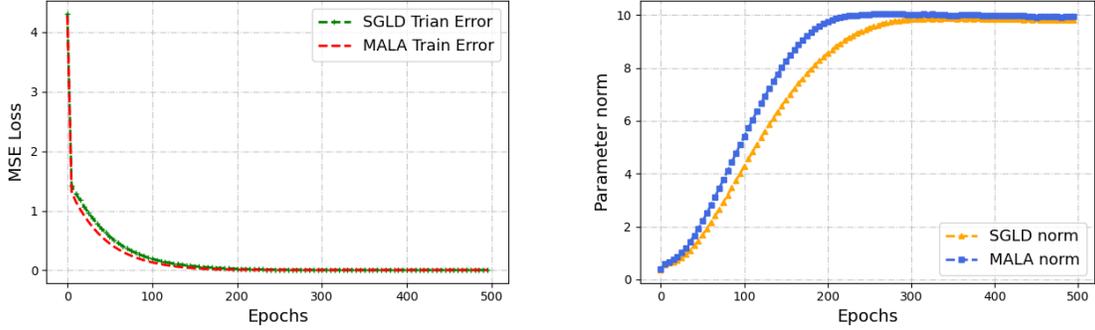


Figure 5: The training MSE losses of MALA and SGLD are presented on the **left**, with a comparative analysis of their  $\ell_2$  norms shown on the **right**. These results were obtained for  $\mathbf{n} = 200$ , using a learning rate of 0.001 and a parameter dimension of  $p = 600$  with  $\lambda = 0.01$ , averaged across ten runs. In terms of computational time, SGLD required an average of 10.4 seconds to converge over 300 epochs, while MALA took an average of 7 seconds to converge over 200 epochs.

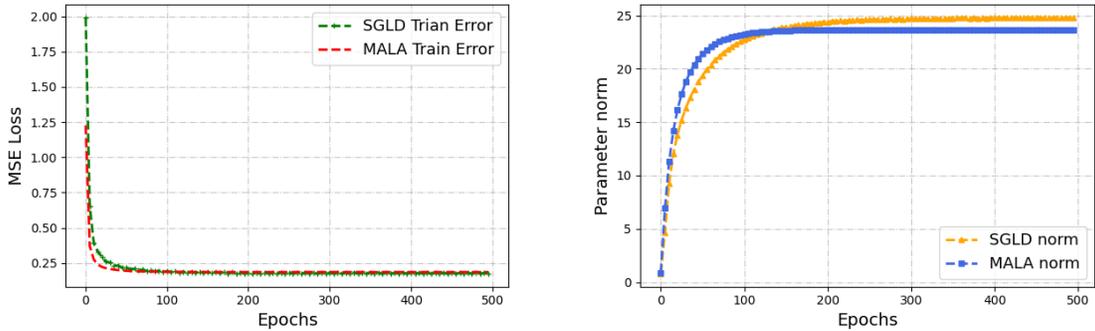


Figure 6: The training MSE losses of MALA and SGLD are presented on the **left**, with a comparative analysis of their  $\ell_2$  norms shown on the **right**. These results were obtained for  $\mathbf{n} = 800$ , using a learning rate of 0.001 and a parameter dimension of  $p = 600$  with  $\lambda = 0.01$ , averaged across ten runs. In terms of computational time, SGLD required an average of 12.4 seconds to converge over 200 epochs, while MALA took an average of 12 seconds to converge over 100 epochs.

the maximum likelihood estimate (MLE), including Assumption 1 and the condition that  $\pi$  is continuous and  $\pi(\mathbf{w}) > 0$  for all  $\mathbf{w} \in \mathcal{W}$ .

For the Gibbs algorithm  $P_{\hat{W}|S}^*$  defined in (10), if we adopt the log-loss function  $\ell(w, \mathbf{z}) = -\log P(y|\mathbf{x}; w)$ , and set  $\beta = n$ , the Gibbs algorithm directly corresponds to the Bayesian posterior of the parametric family with prior distribution  $\pi$ . Therefore, in the asymptotic regime where  $p$  is fixed and  $n \rightarrow \infty$ , Bernstein–von-Mises theorem under model mismatch (Van der Vaart, 2000; Kleijn et al., 2012) states that we could approximate the Gibbs algorithm in (10) by

$$\mathcal{N}(\hat{W}_{\text{ML}}, \frac{1}{n} J(\mathbf{w}^*)^{-1}),$$

with  $\mathbf{w}^*$  and  $J(\mathbf{w})$  defined in a similar manner as in (31). As  $n \rightarrow \infty$ , the asymptotic Normality of the MLE states that the distribution of  $\hat{W}_{\text{ML}}$  will converge to

$$\mathcal{N}(\mathbf{w}^*, \frac{1}{n} J(\mathbf{w}^*)^{-1} I(\mathbf{w}^*) J(\mathbf{w}^*)^{-1}).$$

Thus, the Gibbs distribution  $P_{\hat{W}|S}^*$  can be approximated as a Gaussian channel regardless of the choice of prior  $\pi(\mathbf{w})$ . Then, the symmetrized KL information can be computed using (Palomar and Verdú, 2008, Theorem 14), which characterizes the symmetrized KL information over a vector Gaussian channel, i.e.,

$$\frac{I_{\text{SKL}}(P_{\hat{W}|S}^*, P_S)}{n} = \frac{\text{tr}(I(\mathbf{w}^*) J(\mathbf{w}^*)^{-1})}{n} = \mathcal{O}\left(\frac{p}{n}\right). \quad (46)$$

When the true data distribution is within the model family, we can have  $I(\mathbf{w}^*) = J(\mathbf{w}^*)$ . Therefore, for the Gibbs algorithm  $P_{\hat{W}|S}^*$  defined in (10), if we adopt the log-loss function  $\ell(w, \mathbf{z}) = -\log P(y|\mathbf{x}; w)$ , and set  $\beta = n$ , and let  $n \rightarrow \infty$  with fixed  $p$ , we can get the following asymptotic expression for the generalization error,

$$\frac{I_{\text{SKL}}(P_{\hat{W}|S}^*, P_S)}{\beta} = \frac{p}{n}. \quad (47)$$

Plug this result back to (14), the Gibbs-based AIC can be written as,

$$\text{AIC}^+ = L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \frac{p}{n}. \quad (48)$$

## E PROOF OF PROPOSITION 2

If we adopt the log-loss function  $\ell(w, \mathbf{z}) = -\log P(y|\mathbf{x}; w)$ , and set  $\beta = n$ , the Gibbs distribution in (10) can be viewed as the Bayesian posterior distribution, i.e.,

$$P_{\hat{W}|S}^*(\mathbf{w}|\mathbf{z}^n) = \frac{\pi(\mathbf{w}) \prod_{i=1}^n P(y_i|\mathbf{x}_i; \mathbf{w})}{V(\mathbf{z}^n)}, \quad \text{with } V(\mathbf{z}^n) = \int \pi(\mathbf{w}) \prod_{i=1}^n P(y_i|\mathbf{x}_i; \mathbf{w}) d\mathbf{w}. \quad (49)$$

Therefore,

$$\begin{aligned} & \mathbb{E}_{P_{\hat{W}|S}^*} [L_E(\hat{W}, \mathbf{z}^n)] + \frac{1}{n} D(P_{\hat{W}|S}^* \| \pi) \\ &= \mathbb{E}_{P_{\hat{W}|S}^*} [L_E(\hat{W}, \mathbf{z}^n)] + \frac{1}{n} \mathbb{E}_{P_{\hat{W}|S}^*} \left[ \log \frac{\exp(-nL_E(\hat{W}, \mathbf{z}^n))}{V(\mathbf{z}^n)} \right] \\ &= -\frac{1}{n} \log V(\mathbf{z}^n) \\ &= -\frac{1}{n} \log m(\mathbf{z}^n), \end{aligned} \quad (50)$$

which completes the proof for the first equality.

Using Lemma 3, with the same Bayesian posterior distribution,

$$\begin{aligned} & \mathbb{E}_{\pi \otimes P_S} [L_E(W, S)] - \frac{1}{\beta} D(\pi \| P_{\hat{W}|S}^* | P_S), \\ &= -\frac{1}{\beta} \mathbb{E}_{P_S} [\log \mathbb{E}_{\pi} [e^{-\beta L_E(W, S)}]] \\ &= -\frac{1}{n} \log m(\mathbf{z}^n), \end{aligned} \quad (51)$$

which completes the proof for the second equality.

## F PROOF OF THEOREM 2

Recall the Gibbs-based BIC is given by

$$\text{BIC}^+ \triangleq L_E(\hat{W}_{\text{Gibbs}}, \mathbf{z}^n) + \frac{1}{n} D(P_{\hat{W}|S}^* \| \pi). \quad (52)$$

Thus, we just need to characterize the asymptotic behavior of the KL divergence term in the regime where  $p$  is fixed and  $n \rightarrow \infty$ . Note that The KL divergence can be written as

$$\frac{1}{n} D(P_{\hat{W}|S}^* \| \pi) = -\frac{1}{n} h(P_{\hat{W}|S}^*) - \frac{1}{n} \mathbb{E}_{P_{\hat{W}|S}^*} [\log \pi(\hat{W})], \quad (53)$$

it should be noticed that  $h(P_{\hat{W}|S}^*)$  is the differential entropy, and  $P_{\hat{W}|S}^*$  will converge to  $\mathcal{N}(\hat{W}_{\text{ML}}, \frac{1}{n} J(\mathbf{w}^*)^{-1})$  by Bernstein–von-Mises theorem as shown in Appendix D. Then, the second term above will reduce to

$\frac{1}{n} \log \pi(\hat{W}_{\text{ML}})$ , and converges to zero as  $n \rightarrow \infty$ . Using the same Gaussian approximation, the differential entropy term can be computed as

$$h(P_{\hat{W}|S=z^n}^*) = \frac{p}{2} \log\left(\frac{2\pi e}{n}\right) + \frac{1}{2} \log |J(\mathbf{w}^*)^{-1}|. \quad (54)$$

As  $n \rightarrow \infty$ ,  $\hat{W}_{\text{Gibbs}} \rightarrow \hat{w}_{\text{ML}}$ . Therefore, the Gibbs BIC can be asymptotically approximated by

$$\begin{aligned} \text{BIC}^+ &= L_E(\hat{W}_{\text{Gibbs}}, s) + \frac{1}{n} D(P_{\hat{W}|S=z^n}^* \|\pi) \\ &= L_E(\hat{W}_{\text{Gibbs}}, s) - \frac{1}{n} h(P_{\hat{W}|S=z^n}^*) - \frac{1}{n} \mathbb{E}_{P_{\hat{W}|S=z^n}^*} [\log \pi(\hat{W})] \\ &= L_E(\hat{w}_{\text{ML}}, s) + \frac{p}{2n} \log \frac{n}{2\pi e} + \frac{1}{2n} \log |J(\mathbf{w}^*)| - \frac{1}{n} \log \pi(\hat{W}_{\text{ML}}). \end{aligned} \quad (55)$$

Thus, both the terms  $p \log 2\pi e$ ,  $\log |J(\mathbf{w}^*)|$  and  $\log \pi(\hat{W}_{\text{ML}})$  have order less than  $O(1)$  with respect to the sample size  $n$ , and can be ignored in the Gibbs-based BIC.

Comparing the above result with the log-marginal likelihood in (41), the penalty term in  $\text{BIC}^+$  differs from the classic BIC by  $\frac{p}{2n} \log e$ . This is due to the fact that we evaluate the same marginal likelihood  $m(\mathbf{z}^n)$  using different algorithms, and the empirical risk achieved by SGLD or MALA is different from that of the SGD.

## G GIBBS DISTRIBUTION OF RANDOM FEATURE MODEL

For random feature model with the prior distribution  $\pi(\mathbf{w}) \sim \mathcal{N}(0, \frac{\sigma^2}{\lambda n} \mathbf{I}_p)$  and  $L_E(\mathbf{w}, s) = -\frac{1}{n} \sum_{i=1}^n \log(y_i | \mathbf{x}_i, \mathbf{w})$ , the log-posterior  $\log(P_{\hat{W}|S}^*(\mathbf{w}|S)) \propto \log \pi(\mathbf{w}) + \log(e^{-nL_E(\mathbf{w}, s)})$ , where

$$L_E(\mathbf{w}, s) = \frac{1}{2n\sigma^2} \|\mathbf{Y} - \mathbf{B}\mathbf{w}\|_2^2 + \frac{1}{2} \log(2\pi\sigma^2).$$

Thus, the Gibbs algorithm, in this case, is given by the following Gaussian posterior distribution, as shown in Murphy (2007),

$$P_{\hat{W}|S}^* \sim \mathcal{N}(\hat{W}_\lambda, \Sigma_w), \quad (56)$$

where  $\hat{W}_\lambda = (\lambda n \mathbf{I}_p + \mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{Y}$ , and  $\Sigma_w = \sigma^2 (\lambda n \mathbf{I}_p + \mathbf{B}^\top \mathbf{B})^{-1}$ .

## H PROOF OF THEOREM 3

We note that the KL divergence between two Gaussian distributions can be written as

$$D(\mathcal{N}(\boldsymbol{\mu}_1, \Sigma_1) \|\mathcal{N}(\boldsymbol{\mu}_2, \Sigma_2)) = \frac{1}{2} \left[ (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^\top \Sigma_2^{-1} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) + \text{tr}(\Sigma_2^{-1} \Sigma_1) - p + \log \frac{\det \Sigma_2}{\det \Sigma_1} \right].$$

The KL divergence between the Gibbs posterior of the RF model and the prior can be computed by,

$$\begin{aligned} \frac{1}{n} D(P_{\hat{W}|S=z^n}^* \|\pi) &= \frac{1}{2n} \left[ \hat{W}_\lambda^\top \left( \frac{\sigma^2}{\lambda n} \mathbf{I}_p \right)^{-1} \hat{W}_\lambda + \text{tr} \left( \frac{\lambda n}{\sigma^2} \Sigma_w \right) + \log \frac{\det(\frac{\sigma^2}{\lambda n} \mathbf{I}_p)}{\det(\Sigma_w)} - p \right] \\ &= \frac{1}{2n} \left[ \frac{\lambda n}{\sigma^2} \|\hat{W}_\lambda\|_2^2 + \text{tr} \left( (\mathbf{I}_p + \frac{\mathbf{B}^\top \mathbf{B}}{\lambda n})^{-1} \right) + \log \det \left( \mathbf{I}_p + \frac{\mathbf{B}^\top \mathbf{B}}{\lambda n} \right) - p \right]. \end{aligned} \quad (57)$$

The trace and the log determinant of the random matrix  $\Sigma = \frac{\mathbf{B}^\top \mathbf{B}}{\lambda n} + \mathbf{I}_p$  can be computed using the following results from Pennington and Worah (2017), which characterizes the probability density function of the eigenvalues of the random matrix  $\mathbf{B}^\top \mathbf{B}/n$  in the over-parameterized regime.

**Lemma 4.** (Pennington and Worah, 2017) *Let the matrix  $\mathbf{M} = \frac{1}{n} \mathbf{B}^\top \mathbf{B} \in \mathbb{R}^{p \times p}$ , where  $\mathbf{B} = f\left(\frac{\mathbf{X}\mathbf{F}}{\sqrt{d}}\right) \in \mathbb{R}^{n \times p}$ , all the elements in  $\mathbf{F} \in \mathbb{R}^{d \times p}$  and  $\mathbf{X} \in \mathbb{R}^{n \times d}$  are generated i.i.d from  $\mathcal{N}(0, 1)$ . Suppose that the activation function has zero mean and finite moments, i.e.,*

$$\mathbb{E}[f(\varepsilon)] = 0, \quad \mathbb{E}[f(\varepsilon)^k] < \infty, \quad \text{for } k > 1, \quad \varepsilon \sim \mathcal{N}(0, 1). \quad (58)$$

and define constants  $\eta$  and  $\xi$  as

$$\eta \triangleq \mathbb{E}[f(\varepsilon)^2], \quad \xi \triangleq \mathbb{E}[f'(\varepsilon)]^2, \quad \varepsilon \sim \mathcal{N}(0, 1), \quad (59)$$

as  $n, d, p \rightarrow \infty$  with  $d/p \rightarrow \psi$ ,  $d/n \rightarrow \phi$ , where  $\psi, \phi \in (0, \infty)$ , then the Stieltjes transform  $G(z)$  of the spectral density of random matrix  $\mathbf{M}$  satisfies

$$dF_{\mathbf{M}}(x) = \frac{1}{\pi} \lim_{\epsilon \rightarrow 0^+} \text{Im}G(x - i\epsilon), \quad G(z) = \frac{\psi}{z} A\left(\frac{1}{z\psi}\right) + \frac{1-\psi}{z}, \quad (60)$$

$$A(t) = 1 + (\eta - \xi)tA_\phi(t)A_\psi(t) + \frac{A_\phi(t)A_\psi(t)t\xi}{1 - A_\phi(t)A_\psi(t)t\xi}, \quad (61)$$

where  $A_\phi(t) = 1 + (A(t) - 1)\phi$  and  $A_\psi(t) = 1 + (A(t) - 1)\psi$ .

This lemma characterizes the spectral density of random matrix  $\mathbf{M}$  for any zero-mean activation functions. However, these implicit equations need to be evaluated numerically, and it is hard to obtain a closed-form expression or provide more insights.

If we further assume that the assumptions in (25) are satisfied, i.e.,  $\mathbb{E}[f(\varepsilon)^2] = 1$ , and  $\mathbb{E}[f'(\varepsilon)]^2 = 0$ , then the result in Lemma 4 can be simplified significantly, as the probability density of the eigenvalues for random matrix  $\mathbf{M}$  will converge to the well-known Marchenko-Pastur distribution with shape parameter  $r = p/n$ , i.e.,

$$dF_{\mathbf{M}}(x) \rightarrow \left(1 - \frac{1}{r}\right)^+ \delta(x) + \frac{\sqrt{(x-a)^+(b-x)^+}}{2\pi r x}, \quad (62)$$

as  $n, d, p$  all go to infinity, where  $(z)^+ \triangleq \max\{0, z\}$ , and  $a \triangleq (1 - \sqrt{r})^2$ , and  $b \triangleq (1 + \sqrt{r})^2$ . Thus, we focus on this case to obtain a convenient, closed-form expression for KL divergence.

The following lemma from Sections 2.2.2 and 2.2.3 in Tulino et al. (2004) characterizes the  $\eta$ -transform and Shannon transform of the Marchenko-Pastur distribution.

**Lemma 5.** *The  $\eta$  and Shannon transform of a nonnegative random variable  $X$  are defined as*

$$\eta_X(\gamma) \triangleq \mathbb{E}\left[\frac{1}{1 + \gamma X}\right], \quad \mathcal{V}_X(\gamma) \triangleq \mathbb{E}[\log(1 + \gamma X)], \quad (63)$$

respectively. If  $X$  is distributed according to Marchenko-Pastur distribution with shape parameter  $r = p/n$ , then

$$\eta_X(\gamma) = 1 - \frac{\mathcal{F}(\gamma, r)}{4r\gamma}, \quad (64)$$

$$\mathcal{V}_X(\gamma) = \log\left(1 + \gamma - \frac{1}{4}\mathcal{F}(\gamma, r)\right) + \frac{1}{r} \log\left(1 + \gamma r - \frac{1}{4}\mathcal{F}(\gamma, r)\right) - \frac{1}{4r\gamma}\mathcal{F}(\gamma, r), \quad (65)$$

where

$$\mathcal{F}(\gamma, r) \triangleq \left(\sqrt{\gamma(1 + \sqrt{r})^2 + 1} - \sqrt{\gamma(1 - \sqrt{r})^2 + 1}\right)^2. \quad (66)$$

Equipped with the aforementioned tools from random matrix theory, we could proceed our analysis,

$$\frac{1}{n} \log \det \left(\mathbf{I}_p + \frac{1}{\lambda n} \mathbf{B}^\top \mathbf{B}\right) = \frac{r}{p} \sum_{i=1}^p \log \left(1 + \frac{1}{\lambda} \lambda_i \left(\frac{1}{n} \mathbf{B}^\top \mathbf{B}\right)\right), \quad (67)$$

where the notation  $\lambda_i(\cdot)$  denote the eigenvalues of the matrix for  $i = 1, \dots, p$ . As shown in Lemma 5, we have

$$\frac{1}{p} \sum_{i=1}^p \log \left(1 + \frac{1}{\lambda} \lambda_i \left(\frac{1}{n} \mathbf{B}^\top \mathbf{B}\right)\right) \rightarrow \int_0^\infty \log \left(1 + \frac{x}{\lambda}\right) dF_{\mathbf{M}}^n(x) = \mathcal{V}_X(1/\lambda) \quad (68)$$

almost surely, when  $n, d, p \rightarrow \infty$ ,  $p/n = r$ . Thus, in the over-parameterized regime, we have

$$\frac{1}{n} \log \det \left(\mathbf{I}_p + \frac{1}{\lambda n} \mathbf{B}^\top \mathbf{B}\right) \rightarrow r \cdot \mathcal{V}_X(1/\lambda) = V(1/\lambda, r). \quad (69)$$

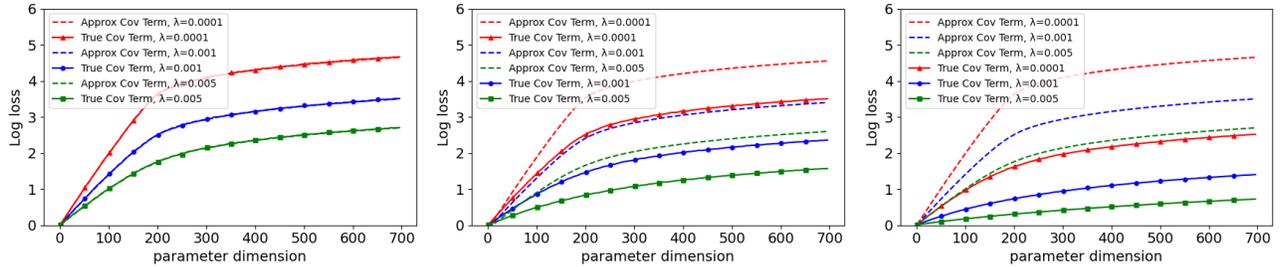


Figure 7: A comparison between  $1/n(\log |\Sigma| + \text{tr}(\Sigma)^{-1} - p)$  and the asymptotic approximation of the covariance term in Theorem 3 for different values of  $\lambda$  and for different activation functions:  $f(x) = (x^2 - 1)/\sqrt{2}$  (**left**), ReLU (**middle**) and Sigmoid (**right**). We adopt the same experiment settings as in Section 5, and we change  $r = p/n$  by fixing  $n = 200$  and varying  $p$ .

The trace term can be simplified as,

$$\frac{1}{n} \text{tr}(\mathbf{I}_p + \frac{1}{\lambda n} \mathbf{B}^\top \mathbf{B})^{-1} = \frac{r}{p} \sum_{i=1}^p \frac{1}{(1 + \frac{1}{\lambda} \lambda_i(\frac{1}{n} \mathbf{B}^\top \mathbf{B}))}, \quad (70)$$

which will converge to the following expression by Lemma 5, when  $n, d, p \rightarrow \infty, p/n = r$ ,

$$\frac{r}{p} \sum_{i=1}^p \frac{1}{(1 + \frac{1}{\lambda} \lambda_i(\frac{1}{n} \mathbf{B}^\top \mathbf{B}))} \rightarrow r \int_0^\infty \frac{1}{1 + \frac{x}{\lambda}} dF_M^n(x) = r(1 - \frac{\mathcal{F}(\frac{1}{\lambda}, r)}{4r\frac{1}{\lambda}}). \quad (71)$$

Combine (69) and (71) with (57), we obtain the following result

$$\frac{1}{n} D(P_{\hat{W}|S=z^n}^* \|\pi) \rightarrow \frac{\lambda}{2\sigma^2} \|\hat{W}_\lambda\|_2^2 - \frac{\lambda}{8} \mathcal{F}(\frac{1}{\lambda}, r) + \frac{1}{2} V(1/\lambda, r). \quad (72)$$

## H.1 Empirical Behavior of Covariance Term

To show that Theorem 3 can provide a good approximation for the asymptotic behavior of the log determinant and trace term in (23), we plot in Figure 7 both the term  $1/n(\log |\Sigma| + \text{tr}(\Sigma)^{-1} - p)$  with finite data and  $\frac{1}{2}V(1/\lambda, r) - \frac{\lambda}{8}\mathcal{F}(\frac{1}{\lambda}, r)$  in the over-parameterized regime for different activation functions with varying regularizer parameters  $\lambda$ . As shown from Figure 7 (left), our theoretical results provide a good proxy for the asymptotic behavior of the covariance term, even for activation functions (e.g., ReLU and Sigmoid in Figure 7 (middle and right)) that do not satisfy the assumptions in Theorem 3. This implies that the particular choice of activation function does not significantly influence the asymptotic behavior of the covariance term in Gibbs-based BIC.

## I ADDITIONAL EXPERIMENTAL RESULTS

### I.1 Experimental Details

In our experiments, the setup for the random feature model is akin to a two-layer neural network with a ReLU activation function. The first layer is initialized with a Gaussian distribution  $\mathcal{N}(0, \frac{1}{\sqrt{d}})$  and remains unchanged throughout the training. The second layer, which starts off with zero values, is trained using the log loss function.

We set the step size for both SGD and MALA to 0.01. The training process of MALA typically converges around 100 epochs. However, to ensure convergence to the Gibbs distribution, we continue to run MALA for a substantial number of epochs (600 in our experiment) even after achieving the minimum training loss. In addition, although we trained the noise variance in the classic regime to calibrate it to an appropriate scale, we fixed the noise variance  $\sigma$  to 0.05 due to the high variability in model complexity in the over-parameterized regime.

Our experiments, implemented in PyTorch requires less than 24 hours of training computation time on a single RTX 3090 GPU. To ensure result accuracy and plot error bars, we perform 50 runs for each setting.

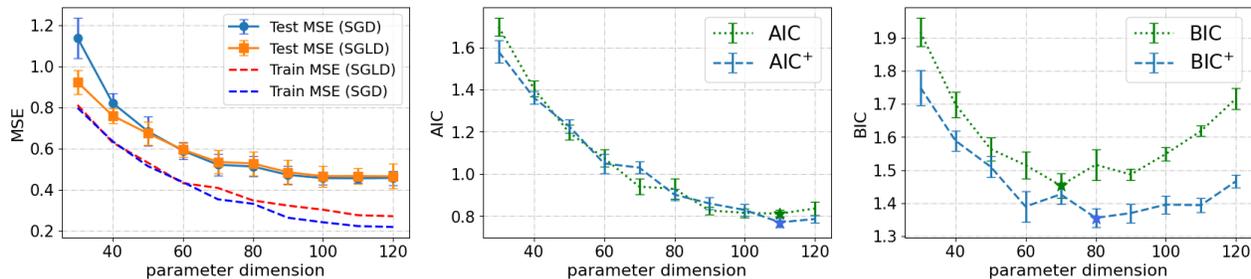


Figure 8: A comparison of SGD and SGLD in terms of MSE (**left**). Comparisons of the classical AIC with AIC<sup>+</sup> in (15) (**middle**), and the classical BIC with BIC<sup>+</sup> in (19) (**right**). All experiments are conducted in the classical  $n \gg p$  setting, with  $n = 600$ ,  $p = [30, \dots, 120]$  and  $\sigma^2 = 0.7$ . The preferred models selected by different information criteria are marked using stars with different colors.

## I.2 RF model loss

Keeping remind that the following loss function is only used in MALA training process,

$$\mathcal{L}(\mathbf{w}) = - \sum_{i=1}^n \left( \log P(y_i | x_i, \mathbf{w}) - \frac{1}{n} \log \pi(\mathbf{w}) \right) \quad (73)$$

$$= \frac{1}{2\sigma^2} \sum_{i=1}^n \left( y_i - f \left( \frac{\mathbf{x}_i^\top \mathbf{F}}{\sqrt{d}} \right) \mathbf{w} \right)^2 + \frac{n}{2} \log(2\pi\sigma^2) + \frac{\lambda n \|\mathbf{w}\|_2^2}{2\sigma^2}. \quad (74)$$

For traditional AIC and BIC, we optimize the following log-likelihood using SGD,

$$\mathcal{L}(\mathbf{w}) = \frac{1}{2\sigma^2} \sum_{i=1}^n \left( y_i - f \left( \frac{\mathbf{x}_i^\top \mathbf{F}}{\sqrt{d}} \right) \mathbf{w} \right)^2 + \frac{n}{2} \log(2\pi\sigma^2). \quad (75)$$

## I.3 Gibbs-based AIC and BIC in Classic Regime

We instantiate a two-layer RF model described in (20), where the first layer is designated for feature mapping and is kept random, and we only train the parameter in the second layer. We use the regularized negative log-likelihood in (22) as the loss function for SGLD to compute both AIC<sup>+</sup> in (15) and BIC<sup>+</sup> in (19), which is different from the SGD algorithm used in classical AIC and BIC.

For simplicity, we generate  $n = 600$  training samples from this linear model

$$y_i = \mathbf{x}_i^\top \mathbf{w}^* + \epsilon_i, \quad \mathbf{w}^* \in \mathbb{R}^p, \quad \|\mathbf{w}^*\|_2^2 = 1, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_d^2), \quad (76)$$

with  $p = 80$ , and noise  $\sigma_d^2 = 0.2$ .

In Figure 8 (Left), we plot the training and test MSEs achieved using SGD and SGLD, respectively. Although SGD and SGLD might yield different model parameters, they exhibit a similar trend in terms of MSE and Log-Loss. The discrepancy in the training algorithms does not significantly impact model selection. Figure 8 (middle) demonstrates that both AIC and AIC<sup>+</sup> follow a similar trend to the test MSE, selecting the model ( $p = 110$ ) that achieves the smallest population risk. On the other hand, due to larger penalty terms, BIC and BIC<sup>+</sup> favor simpler models and identify the model most likely to generate the training data. As shown in Figure 8 (right), the classical BIC selects the model with  $p = 70$ , while BIC<sup>+</sup> favors the model with  $p = 80$ . In summary, our observations suggest that AIC<sup>+</sup> and BIC<sup>+</sup> based on SGLD exhibit similar model selection capabilities to the traditional AIC and BIC in the classical large  $n$  regime.

## I.4 Additional Baseline comparison

In Figure 9, we generated  $n = 200$  data by a RF model  $y_i = f \left( \frac{\mathbf{x}_i^\top \mathbf{F}}{\sqrt{d}} \right) \mathbf{w} + \epsilon_i$ , with  $F \in \mathbb{R}^{400 \times 600}$ , i.e., the number of random features  $p = 600$ , and the dimension of  $X$  is  $d = 400$ . The variance of the noise  $\epsilon_i \sim \mathcal{N}(0, 0.1)$ , which

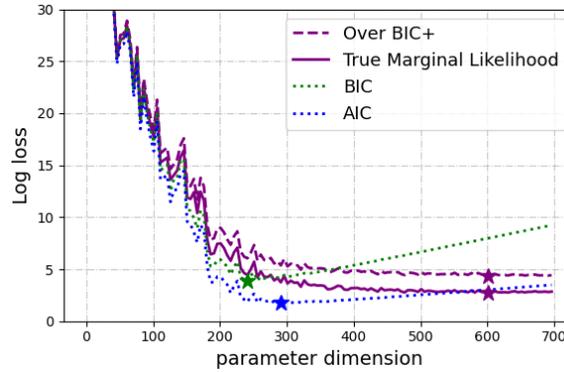


Figure 9: the comparison of the proposed over-parameterized  $BIC^+$  with true marginal likelihood and other classical information criteria for over-parameterized RF models plotted with varying parameter dimension  $p$ .

is different from the  $\sigma = 0.2$  in the likelihood function and the Gaussian prior. We fit the data using the RF model with different dimensions  $p$  to see if our  $BIC^+$  can effectively identify the correct model order, rather than just selecting larger models.

We randomly select  $p$  features from the true  $F$  when  $p \leq 600$ , and we generate extra random features using standard Gaussian if  $p$  goes beyond 600. As demonstrated in Figure 9, our Gibbs-based  $BIC^+$  can reliably select the true model with moderate size instead of simply picking the largest model. In addition, we also plot the baseline that represents the true marginal likelihood without any approximation. Our  $BIC^+$  has a similar behavior, and both  $BIC^+$  and the true marginal likelihood select the correct model with  $p = 600$ .