

EXCESS RISK ANALYSIS FOR EPISTEMIC UNCERTAINTY WITH APPLICATION TO VARIATIONAL INFERENCE

Anonymous authors

Paper under double-blind review

ABSTRACT

Bayesian deep learning plays an important role especially for its ability evaluating epistemic uncertainty (EU). Due to computational complexity issues, approximation methods such as variational inference (VI) have been used in practice to obtain posterior distributions and their generalization abilities have been analyzed extensively, for example, by PAC-Bayesian theory; however, little analysis exists on EU, although many numerical experiments have been conducted on it. In this study, we analyze the EU of supervised learning in approximate Bayesian inference by focusing on its excess risk. First, we theoretically show the novel relations between generalization error and the widely used EU measurements, such as the variance and mutual information of predictive distribution, and derive their convergence behaviors. Next, we clarify how the objective function of VI regularizes the EU. With this analysis, we propose a new objective function for VI that directly controls the prediction performance and the EU based on the PAC-Bayesian theory. Numerical experiments show that our algorithm significantly improves the EU evaluation over the existing VI methods.

1 INTRODUCTION

As machine learning applications spread, understanding the uncertainty of predictions is becoming more important to increase our confidence in machine learning algorithms (Bhatt et al., 2021). Uncertainty refers to the variability of a prediction caused by missing information. For example, in regression problems, it corresponds to the error bars in predictions; and in classification problems, it is often expressed as the class posterior probability, entropy, and mutual information (Hüllermeier & Waegeman, 2021; Gawlikowski et al., 2022). There are two types of uncertainty (Bhatt et al., 2021): 1) Aleatoric uncertainty (AU), which is caused by noise in the data itself, and 2) Epistemic uncertainty (EU), which is caused by a lack of training data. In particular, since EU can tell us where in the input space is yet to be learned, integrated with deep learning methods, it is used in such applications as dataset shift (Ovadia et al., 2019), adversarial data detection (Ye & Zhu, 2018), active learning (Houlsby et al., 2011), Bayesian optimization (Hernández-Lobato et al., 2014), and reinforcement learning (Janz et al., 2019).

Mathematically, AU is defined as Bayes risk, which expresses the fundamental difficulty of learning problems (Depeweg et al., 2018; Jain et al., 2021; Xu, 2020). For EU, Bayesian inference is useful because posterior distribution updated from prior distribution can represent a lack of data (Hüllermeier & Waegeman, 2021). In practice, measurements like the variance of the posterior predictive distribution, and associated conditional mutual information represented EU in practice (Kendall & Gal, 2017; Depeweg et al., 2018).

In Bayesian inference, since posterior distribution is characterized by the training data and the model using Bayes' formula, its prediction performance and EU are determined automatically. However, due to computational issues, such exact Bayesian inference is difficult to implement; we often use approximation methods, such as variational inference (VI) (Bishop, 2006), especially for deep Bayesian models. Since the derived posterior distribution also depends on the properties of approximation methods, the prediction performance and EU of deep Bayesian learning are no longer automatically guaranteed through Bayes' formula. The prediction performance has been analyzed as generalization error, for example, by PAC-Bayesian theory (Alquier, 2021). Since EU is also essential

in practice, we must obtain a theoretical guarantee of the algorithm- and the sample- dependent non-asymptotic theory for EU, similarly to generalization error analysis.

Unfortunately, study has been limited in that direction. Traditional EU analysis has focused on the properties of the exact Bayesian posterior and predictive distributions (Fiedler et al., 2021; Lederer et al., 2019) as well as large sample behaviors (Clarke & Barron, 1990). Since Bayesian deep learning uses approximate posterior distributions, we cannot apply such traditional EU analysis based on Bayes’ formula to Bayesian deep learning. The asymptotic theory of a sufficiently large sample may overlook an important property of the EU that is due to the lack of training data.

Recently, analysis of EU focusing on loss functions was proposed for supervised learning (Xu & Raginsky, 2020; Jain et al., 2021). EU was defined as the **excess risk** obtained by subtracting the Bayes risk corresponding to the AU from the total risk. Thus, excess risk implies the loss due to insufficient data when the model is well specified. Although this approach successfully defines EU with loss functions, the following limitation still remains. Xu & Raginsky (2020) assume that the data generating mechanism is already known and that we can precisely evaluate Bayesian posterior and predictive distribution. A correct model is not necessarily a realistic assumption, and the assumption about an exact Bayesian posterior hampers understanding EU in approximation methods.

To address these limitations, it appears reasonable to analyze excess risk under a similar setting as PAC-Bayesian theory and apply it to EU of approximate Bayesian inference. However, as shown in Sec. 2, analyzing excess risk in such a way leads to impractical theoretical results, and the relations between excess risk and the widely used EU measurements remain unclear. This greatly complicates EU analysis. Because of this difficulty, to the best of our knowledge, no research exists on excess risk for EU for approximation methods.

In this paper, we propose a new theoretical analysis for EU that addresses the above limitations of these existing settings. Our contributions are the followings:

- We show non-asymptotic analysis for widely used EU measurements (Theorems 2 and 3). We propose computing the Bayesian excess risk (BER) (Eq. (9)) and show that this excess risk equals to widely used EU measurements. Then we theoretically show the convergence behavior of BER using PAC-Bayesian theory (Eqs. (13)) and (18)).
- Based on theoretical analysis, we give a new interpretation of the existing VI that clarifies how the EU is regularized (Eqs. (19) and (20)). Then we propose a novel algorithm that directly controls the prediction and the EU estimation performance simultaneously based on PAC-Bayesian theory (Eq. (21)). Numerical experiments suggest that our algorithm significantly improves EU evaluation over the existing VI.

2 BACKGROUND OF PAC-BAYESIAN THEORY AND EPISTEMIC UNCERTAINTY

Here we introduce preliminaries. Such capital letters as X represent random variables, and such lowercase letters as x represent deterministic values. All the notations are summarized in Appendix A. In Appendix B, we show summary of the settings.

2.1 PAC-BAYESIAN THEORY

We consider a supervised setting and denote input-output pairs by $Z = (X, Y) \in \mathcal{Z} := \mathcal{X} \times \mathcal{Y}$. We assume that all the data are i.i.d. from some unknown data-generating distribution $\nu(Z) = \nu(Y|X)\nu(X)$. Learners can access N training data, $\mathbf{Z}^N := (Z_1, \dots, Z_N)$ with $Z_n := (X_n, Y_n)$, which are generated by $\mathbf{Z}^N \sim \nu(Z)^N$. We express $\nu(Z)^N$ as $\nu(\mathbf{Z}^N)$. We express conditional distribution as $\nu(Y|X = x)$ as $\nu(Y|x)$ for simplicity. We introduce loss function $l : \mathcal{Y} \times \mathcal{A} \rightarrow \mathbb{R}$ where \mathcal{A} is an action space. We express loss of action $a \in \mathcal{A}$ and target variable y is written as $l(y, a)$. We introduce a model $f_\theta : \mathcal{X} \rightarrow \mathcal{A}$, parameterized by $\theta \in \Theta \subset \mathbb{R}^d$. When we put a prior $p(\theta)$ over θ , the PAC-Bayesian theory (Alquier, 2021; Germain et al., 2016) guarantees the prediction performance by focusing on the average of the loss with respect to posterior distribution $q(\theta|\mathbf{Z}^N) \in \mathcal{Q}$. \mathcal{Q} is a family of distributions and $q(\theta|\mathbf{Z}^N)$ is not restricted to Bayesian posterior distribution. In this work we consider the log loss and the squared loss. For the log loss, we consider model $p(y|x, \theta)$, and the loss is given as $l(y, p(y|x, \theta)) = -\ln p(y|x, \theta)$, where \mathcal{A} is probability distributions. For the squared loss, we use model $f_\theta(x)$ and $l(y, f_\theta(x)) = |y - f_\theta(x)|^2$, where $\mathcal{Y} = \mathcal{A} = \mathbb{R}$.

PAC-Bayesian theory provides a guarantee for the generalization of test error $R^l(Y|X, \mathbf{Z}^N) := \mathbb{E}_{\nu(\mathbf{Z}^N)} \mathbb{E}_{q(\theta|\mathbf{Z}^N)} \mathbb{E}_{\nu(Z)} l(Y, f_\theta(X))$ and training error $\mathbb{E}_{\nu(\mathbf{Z}^N)} r^l(\mathbf{Z}^N) := \mathbb{E}_{\nu(\mathbf{Z}^N)} \mathbb{E}_{q(\theta|\mathbf{Z}^N)} \frac{1}{N} \sum_{n=1}^N l(Y_n, f_\theta(X_n))$. A typical PAC-Bayesian error bound takes form $R^l(Y|X, \mathbf{Z}^N) \leq \mathbb{E}_{\nu(\mathbf{Z}^N)} r^l(\mathbf{Z}^N) + \text{Gen}^l(\mathbf{Z}^N)$. $\text{Gen}^l(\mathbf{Z}^N)$ is called **generalization error**. Under suitable assumptions (Alquier, 2021), $\text{Gen}^l(\mathbf{Z}^N)$ is upper-bounded by $\mathcal{O}(1/N^\alpha)$ for $\alpha \in (1/2, 1]$. In many cases, it depends on the complexity of the posterior distribution, such as Kullback-Leibler (KL) divergence $\text{KL}(q(\theta|\mathbf{Z}^N)|p(\theta))$. When $\text{Gen}^l(\mathbf{Z}^N) = \text{KL}(q(\theta|\mathbf{Z}^N)|p(\theta))/\lambda + c$, where λ and c are positive constants, given training data $\mathbf{Z}^N = \mathbf{z}^N$, we get a posterior distribution for the prediction by

$$\hat{q}(\theta|\mathbf{z}^N) = \arg \min_{q(\theta|\mathbf{z}^N) \in \mathcal{Q}} r(\mathbf{z}^N) + \frac{\text{KL}(q(\theta|\mathbf{z}^N)|p(\theta))}{\lambda}. \quad (1)$$

When the log loss and $\lambda = N$ is used, this minimization is closely related to variational inference (VI) in Bayesian inference. See (Germain et al., 2016) for details.

Under additional moderate assumptions, using $\hat{q}(\theta|\mathbf{z}^N)$ for the test error, we can derive the following **excess risk (ER)** bound from the PAC-Bayesian generalization bound (Alquier, 2021):

$$\text{ER}^l(Y|X, \mathbf{Z}^N, \theta^*) := R^l(Y|X, \mathbf{Z}^N) - R^l(Y|X, \theta^*) \leq C_1 \frac{\ln N}{N^\alpha}, \quad (2)$$

where $R^l(Y|X, \theta^*) = \mathbb{E}_{\nu(Z)} l(Y, f_{\theta^*}(X))$ and $\theta^* = \arg \min_{\theta} \mathbb{E}_{\nu(Z)} l(Y, f_\theta(X))$. Constant C_1 depends only on the problem. Since we aim to analyze EU, we do not further discuss the details of the PAC-Bayesian bound. See Appendix C.2 for the explicit conditions of this bound.

Although PAC-Bayesian theory focuses on the average test error over posterior distribution, we use predictive distribution for predictions in Bayesian inference. Thus we define **Prediction Risk (PR)**:

$$\text{PR}^l(Y|X, \mathbf{Z}^N) := \mathbb{E}_{\nu(\mathbf{Z}^N)} \mathbb{E}_{\nu(Z)} l(Y, \mathbb{E}_{q(\theta|\mathbf{Z}^N)} f_\theta(X)). \quad (3)$$

When the loss is log loss, $\text{PR}^{\log}(Y|X, \mathbf{Z}^N) = -\mathbb{E}_{\nu(\mathbf{Z}^N)} \mathbb{E}_{\nu(Z)} \log p^q(Y|X, \mathbf{Z}^N)$ where $p^q(y|x, \mathbf{z}^N) := \mathbb{E}_{q(\theta|\mathbf{z}^N)} p(y|x, \theta)$ is the approximate predictive distribution. Thus, $\text{PR}^{\log}(Y|X, \mathbf{Z}^N)$ corresponds to the log loss of the predictive distribution, which is commonly used in the analysis of Bayesian inference (Watanabe, 2009; 2018).

2.2 EPISTEMIC UNCERTAINTY MEASUREMENTS

Here, we introduce widely used EU measurements in approximate Bayesian inference. For the log loss, conditioned on $(X, \mathbf{Z}^N) = (x, \mathbf{z}^N)$, the approximate mutual information has been widely used for uncertainty estimation (Depeweg et al., 2018):

$$I_\nu(\theta; Y|x, \mathbf{z}^N) = H[p^q(Y|x, \mathbf{z}^N)] - \mathbb{E}_{q(\theta|\mathbf{z}^N)} H[p(Y|x, \theta)], \quad (4)$$

where $H[p^q(Y|x, \mathbf{z}^N)] := -\mathbb{E}_{p^q(Y|x, \mathbf{z}^N)} \log p^q(Y|x, \mathbf{z}^N)$ is the entropy of the approximate predictive distribution and $\mathbb{E}_{q(\theta|\mathbf{z}^N)} H[p(Y|x, \theta)]$ is the conditional entropy. $I_\nu(\theta; Y|x, \mathbf{z}^N)$ has been used in Bayesian experimental design (Foster et al., 2019) and reinforcement learning (Depeweg et al., 2018). Note that by taking the expectation, we have $I_\nu(\theta; Y|X, \mathbf{Z}^N) = \mathbb{E}_{\nu(X=x)} I_\nu(\mathbf{Z}^N=\mathbf{z}^N)(\theta; Y|x, \mathbf{z}^N)$.

In the case of squared loss, the variance of the model is often used for EU. This is a common practice in VI, Monte Carlo (MC) dropout (Kendall & Gal, 2017), and deep ensemble methods (Lakshminarayanan et al., 2017). Conditioned on $(X, \mathbf{Z}^N) = (x, \mathbf{z}^N)$, it is written as

$$\text{Var}_{\theta|\mathbf{z}^N} f_\theta(x) = \mathbb{E}_{q(\theta|\mathbf{z}^N)} (f_\theta(x) - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(x))^2. \quad (5)$$

Although Eqs. (4) and (5) are widely used in application, there have been limited theoretical study for them as discussed in Sec. 1.

2.3 EXCESS RISK ANALYSIS AND EPISTEMIC UNCERTAINTY

Recently, the analysis of EU based on excess risk was proposed (Xu & Raginsky, 2020). The key idea of this analysis is to assume that our statistical model $p(y|x, \theta)$ is correct and address the average performance of this model by assuming a prior distribution over θ with distribution $p(\theta)$.

Specifically, the joint distribution of the training data, the test data, and parameter of the model is given as $p_B(\mathbf{Z}^N, Z, \theta) := p(\theta) \prod_{n=1}^N p(Y_n|X_n, \theta)\nu(X_n)p(Y|X, \theta)\nu(X)$. Under this setting, the goal of learning is to infer decision rule $\psi : \mathcal{Z}^N \times \mathcal{X} \rightarrow \mathcal{A}$ that minimizes expected loss $\mathbb{E}_{p_B(\mathbf{Z}^N, Z, \theta)}[l(Y, \psi(X, \mathbf{Z}^N))]$. They refer to this setting as Bayesian learning since we marginalize out parameter θ . With this notation, Xu & Raginsky (2020) defined **minimum excess risk** as

$$\text{MER}^l(Y|X, \mathbf{Z}^N) := \inf_{\psi: \mathcal{Z}^N \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p_B(\mathbf{Z}^N, Z, \theta)}[l(Y, \psi(X, \mathbf{Z}^N))] - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p_B(\mathbf{Z}^N, Z, \theta)}[l(Y, \phi(\theta, X))], \quad (6)$$

where the first term is the minimum achievable risk using the training data and the second term is the Bayes risk since it uses learning rule $\phi : \Theta \times \mathcal{X} \rightarrow \mathcal{A}$, which takes true parameter θ instead of the training data. Thus, the second term is the aleatronic uncertainty (AU) since it expresses the task's fundamental difficulty. Then MER can be regarded as the EU since it is the difference between the total risk and the AU (Xu & Raginsky, 2020; Hafez-Kolahi et al., 2021).

For the log loss, the first term is $H[p(Y|X, \mathbf{Z}^N)]$ and the second term is $\mathbb{E}_{p(\theta)}H[p(Y|X, \theta)]$. Thus, $\text{MER}^{\text{log}}(Y|X, \mathbf{Z}^N) = I(\theta; Y|X, \mathbf{Z}^N)$, which is the conditional mutual information. Other than the log loss, if the loss function satisfies the σ^2 sub-Gaussian property conditioned on $(X, \mathbf{Z}^N) = (x, \mathbf{z}^N)$, $\text{MER}^l(Y|X, \mathbf{Z}^N) \leq \sqrt{2\sigma^2}I(\theta; Y|X, \mathbf{Z}^N)$ holds (Xu & Raginsky, 2020). In many practical settings $I(\theta; Y|X, \mathbf{Z}^N)$ is upper-bounded by $\mathcal{O}(\ln N/N)$. Thus, EU converges with $\mathcal{O}(\ln N/N)$ under this settings. See Appendix C for more details.

Although this analysis successfully defined EU with rigorous theoretical analysis, the assumptions are clearly impractical since we assume that the correct model, exact Bayesian posterior, and predictive distributions are available. To extend this analysis into approximate Bayesian inference, it is tempting to combine the theory of MER with PAC-Bayesian theory where the data are generated i.i.d from $\nu(Z)$. For that extension, here we introduce the **Prediction Excess Risk (PER)** using Eq. (3):

$$\text{PER}^l(Y|X, \mathbf{Z}^N) := \text{PR}^l(Y|X, \mathbf{Z}^N) - \inf_{\tilde{\phi}: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Z)}[l(Y, \tilde{\phi}(X))], \quad (7)$$

where the second term corresponds to the Bayes risk. Although we introduced this definition inspired by Eq. (6), it is impractical for evaluating EU. In practice, we are interested in evaluating EU using only input x , as shown in Eqs. (4) and (5). However, we cannot use Eq. (7) for that purpose since we do not know both $\nu(Y|x)$ and the Bayes risk in the second term. Despite less practical definition, as shown in Sec.3, PER plays a fundamental role in understanding the algorithm-dependent behavior of the widely used EU measurements in Eqs. (4) and (5).

3 ANALYSIS OF EPISTEMIC UNCERTAINTY BASED ON EXCESS RISK

In this section, we develop theories for analyzing the widely used EU measurements introduced in Sec 2.2. We focus on the following questions. **(Q1)** The convergence behaviors of those measurements are not apparent. As the number of training data points increases, we expect these measurements to converge to zero. **(Q2)** The relationship between these measurements and the generalization error is unclear. Since these measurements depend on the training data and the algorithm, we expect some meaningful relationships must exist. All the proofs in this section are shown in Appendix D.

3.1 RELATION BETWEEN EPISTEMIC UNCERTAINTY AND BAYESIAN EXCESS RISK

First, to connect the practical EU evaluation methods in Sec.2.2 with the excess risk analysis in Sec.2.3, we introduce the approximate joint distribution of test data, training data, and parameters:

$$\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(Z) \approx p^q(\theta, \mathbf{Z}^N, Z) := \nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X)p(Y|X, \theta). \quad (8)$$

When the log loss is used, we employ model for $p(y|x, \theta)$ in Eq. (8). When the squared loss is used, we assume Gaussian distribution $p(y|x, \theta) = N(y|f_\theta(x), v^2)$ for some $v \in \mathbb{R}$.

When a model is well specified, that is, $\nu(y|x) = p(y|x, \theta^*)$ holds for some $\theta^* \in \Theta$, we expect that the predictive distribution converges to $p(y|x, \theta^*)$ and the approximation of Eq. (8) becomes accurate as N increases. We discuss the quality of this approximation in Sec.3.2. Under this setting, we define **Bayesian Excess risk (BER)**:

$$\text{BER}^l(Y|X, \mathbf{Z}^N) := \text{BPR}^l(Y|X, \mathbf{Z}^N) - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)}l(Y, \phi(\theta, X)) \quad (9)$$

where BPR is the **Bayesian Prediction Risk** defined as

$$\text{BPR}^l(Y|X, \mathbf{Z}^N) := \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)} l(Y, \mathbb{E}_{q(\theta|\mathbf{Z}^N)} f_{\theta'}(X)), \quad (10)$$

and the second term is the Bayes risk under the approximate joint distribution of Eq. (8). Note that $\text{BER}^l(Y|X, \mathbf{Z}^N)$ is always larger than 0; see Appendix D.2 for details. We also show the formal definitions of BER and BPR conditioned on $(X, \mathbf{Z}^N) = (x, \mathbf{z}^N)$ in Appendix D.1.

BER and BPR are defined, motivated by PER, PR, and MER. The difference is the mechanism of through which the test data are generated. In BER, we assume that our model $p(y|x, \theta)$ is correct, and the parameters follow the approximate posterior distribution $q(\theta|\mathbf{z}^N)$. Thus, the data-generating mechanism resembles the setting in Sec.2.3. Therefore, similar to MER, BER implies the loss due to insufficient data under the assumption that our current model $q(\theta|\mathbf{z}^N)p(y|x, \theta)$ is correct.

The next theorem elaborate this intuition and connects BER to widely used EU measurements:

Theorem 1. *Conditioned on (x, \mathbf{z}^N) , we express $\text{BER}^{\log}(Y|x, \mathbf{z}^N)$ for the log loss and $\text{BER}^{(2)}(Y|x, \mathbf{z}^N)$ for the squared loss. Under the definition of Eq. (9), we have*

$$\text{BER}^{\log}(Y|x, \mathbf{z}^N) = I_{\nu}(\theta; Y|x, \mathbf{z}^N), \quad \text{BER}^{(2)}(Y|x, \mathbf{z}^N) = \text{Var}_{\theta|\mathbf{z}^N} f_{\theta}(x). \quad (11)$$

Thus, by studying BER, we can analyze the widely used EU measurements. We point out that $\text{BPR}^{\log}(Y|x, \mathbf{z}^N) = H[p^q(Y|x, \mathbf{z}^N)]$, and the Bayes risk is given as $\mathbb{E}_{q(\theta|\mathbf{z}^N)} H[p(Y|x, \theta)]$.

Remark 1. *BER captures EU when our model is well specified. On the other hand, PER (Eq. (7)) represents the prediction performance, which has the relation to the quality of approximation of Eq. (8) under the given loss function. This intuition leads to our new VI in Sec.3.3.*

3.2 ANALYSIS OF EXCESS RISK BASED ON PAC-BAYESIAN THEORY

Based on the definitions introduced in Sec. 3.1, here we develop a novel relation between BER and the generalization. First, we show the results of the squared loss. For simplicity, assume $\mathcal{Y} = \mathbb{R}$. See Appendix D.6 for $\mathcal{Y} = \mathbb{R}^d$.

Theorem 2. *Conditioned on (x, \mathbf{z}^N) , assume that a regression function is well specified, that is, $\mathbb{E}_{\nu(Y|x)}[Y|x] = f_{\theta^*}(x)$ holds. Then we have*

$$\text{PER}^{(2)}(Y|x, \mathbf{z}^N) + \text{BER}^{(2)}(Y|x, \mathbf{z}^N) = \text{ER}^{(2)}(Y|x, \mathbf{z}^N, \theta^*) \leq R^{(2)}(Y|x, \mathbf{z}^N). \quad (12)$$

Furthermore assume that the PAC-Bayesian bound Eq. (2) holds, and then we have

$$\text{PER}^{(2)}(Y|X, \mathbf{Z}^N) + \text{BER}^{(2)}(Y|X, \mathbf{Z}^N) = \text{ER}^{(2)}(Y|X, \mathbf{Z}^N, \theta^*) \leq \frac{C_1 \ln N}{N^\alpha}. \quad (13)$$

Proof. We use the following relation about the Jensen gap and BER:

Lemma 1. *For any (x, y) and any posterior distribution conditioned on $\mathbf{Z}^N = \mathbf{z}^N$, we have*

$$|y - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_{\theta}(x)|^2 + \text{Var}_{\theta|\mathbf{z}^N} f_{\theta}(x) = \mathbb{E}_{q(\theta|\mathbf{z}^N)} |y - f_{\theta}(x)|^2. \quad (14)$$

From this lemma, the theorem follows directly. \square

Remark 2. *When we use a flexible model, such as a deep neural network for $f_{\theta}(x)$, assumption $\mathbb{E}_{\nu(Y|x)}[Y|x] = f_{\theta^*}(x)$ holds even when $\nu(Y|x) \neq p(y|x, \theta^*)$, which means we misspecify the noise function in regression tasks.*

From Eq. (12), $\text{Var}_{\theta|\mathbf{z}^N} f_{\theta}(x)$ clearly is a lower bound of excess risk and test error, consistent with the well-known result that the variance of the predictor often underestimates EU (Lakshminarayanan et al., 2017). From Eq. (13), $\text{BER}^{(2)}(Y|X, \mathbf{Z}^N)$ converges to 0 with the same order as the PAC-Bayesian bound. Finally, we remark that from Lemma 1, we have

$$R^{(2)}(Y|X, \mathbf{Z}^N) = \text{PR}^{(2)}(Y|X, \mathbf{Z}^N) + \text{BER}^{(2)}(Y|X, \mathbf{Z}^N). \quad (15)$$

This indicates that the test error is decomposed into PR and BER. As pointed out in Remark 1, BER is EU under the approximation of Eq. (8), and PER represents the quality of that approximation, Eq. (15) suggests that the test error simultaneously regularizes those BER and PER.

Next we show the log loss result. Our analysis requires additional assumption about model $p(y|x, \theta)$. We define log density ratio $L(y, x, \theta, \theta^*) := -\ln p(y|x, \theta) + \ln p(y|x, \theta^*)$.

Assumption 1. Conditioned on $(x, \theta, \mathbf{z}^N)$, there exists convex function $h(\rho)$ for $[0, b)$ such that cumulant function $L(y, x, \theta, \theta^*)$ is upper-bounded by $h(\rho)$, i.e., the following inequality holds:

$$\ln \mathbb{E}_{p(Y|x, \theta)} e^{\rho(L(Y, x, \theta, \theta^*) - \mathbb{E}_{p(Y|x, \theta)} L(Y, x, \theta, \theta^*))} \leq h(\rho). \quad (16)$$

For example, if $h(\rho) = \rho^2 \sigma^2(x, \theta)/2$ and $b = \infty$, this assumption resembles the σ^2 sub-Gaussian property given $(x, \theta, \mathbf{z}^N)$. When considering Gaussian likelihood $p(y|x, \theta) = N(y|f_\theta(x), v^2)$, we have $h(\rho) = \frac{\rho^2}{2v^2} |f_\theta(x) - f_{\theta^*}(x)|^2$. Thus, $\sigma^2(x, \theta)$ depends on x and θ , and we refer to this $\sigma^2(x, \theta)$ as a sub-Gaussian property. Other than the Gaussian likelihood, when the log loss is bounded, it satisfies the sub-Gaussian property. In this paper, we focus on this sub-Gaussian setting for Assumption 1 to clarify the presentation. We show an example of the logistic regression in Appendix D.11.

Theorem 3. When the model is well specified, that is, $\nu(y|x) = p(y|x, \theta^*)$ holds and $\sigma^2(x, \theta)$ sub-Gaussian property is satisfied for $L(y, x, \theta, \theta^*)$, as discussed above. Assume that $\mathbb{E}_{q(\theta|\mathbf{z}^N)} \sigma^2(x, \theta) < \sigma_p^2 < \infty$. Conditioned on (x, \mathbf{z}^N) , we have

$$\text{PER}^{\log}(Y|x, \mathbf{z}^N) + \text{BER}^{\log}(Y|x, \mathbf{z}^N) \leq \sqrt{2\sigma_p^2 \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*)}. \quad (17)$$

Moreover, assume PAC-Bayesian bound Eq. (2) and $\mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X)} \sigma^2(X, \theta) < \sigma_q^2 < \infty$ hold, and then we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq \sqrt{2\sigma_q^2 \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*)} \leq \sqrt{\frac{2\sigma_q^2 C_1 \ln N}{N^\alpha}}. \quad (18)$$

Proof. Using Assumption 1, we apply change-of-measure inequalities (Ohnishi & Honorio, 2021) to control the approximation error of Eq. (8). \square

Remark 3. When we use a generalized linear model, the assumption of a well-specified model is relaxed so that $\mathbb{E}_{\nu(Y|x)}[Y|x]$ is well specified, similar to Remark 2. See Appendix D.9 for details.

From Eq. (17), $I_\nu(\theta; Y|x, \mathbf{z}^N)$ is a lower bound of the excess risk and test error. From Eq. (18), $I_\nu(\theta; Y|X, \mathbf{Z}^N)$ converges in the order of $\mathcal{O}(\sqrt{\ln N/N^\alpha})$ if we can upper-bound $\mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X)} \sigma^2(X, \theta) < \sigma_q^2 < \infty$. For the Gaussian likelihood, we have $\mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X)} \sigma^2(X, \theta) \leq 2\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*) \leq \frac{2C_1 \ln N}{N^\alpha} := \sigma_q^2$. See Appendix D.8 for a detail.

In a similar way, we can derive the convergence rate of the entropy of the predictive distribution, which shows $H[p^q(Y|X, \mathbf{Z}^N)] = H[p(Y|X, \theta^*)] + \mathcal{O}(\sqrt{\ln N/N^\alpha})$. See Appendix D.10 for a formal statement. We show similar results for Theorem 3 under sub-exponential property in Appendix E.

In summary, we obtained the convergence of widely used EU measurements and the entropy in the approximate Bayesian inference for the first time. They converges faster than excess risks. Moreover we obtained two messages from Theorems 2 and 3. First, the widely used EU measurements are the lower bounds of the test error and excess risks. This is consistent with the experimental fact that these EU measurements often underestimate EU. Second, the sum of PER and BER is upper-bounded by excess risk (test error). Thus, when minimizing the test error, we also simultaneously minimize PER and BER. This interpretation extends the intuition of Remark 1 and leads to a new VI in Sec. 3.3.

3.3 NOVEL EU REGULARIZATION METHOD FOR VARIATIONAL INFERENCE

As seen in Sec. 3.2, minimizing the test error leads to minimizing BER and PER. In this section, we discuss this relation using the objective function of VI. As an explicit example, consider a regression problem using $N(y|f_\theta(x), v^2)$. Then from Lemma 1, the loss function of the standard VI (eliminating $\text{KL}(q(\theta|\mathbf{z}^N)|p(\theta))$) can be written:

$$\begin{aligned} -\mathbb{E}_{\nu(Z)} \mathbb{E}_{q(\theta|\mathbf{z}^N)} \ln N(y|f_\theta(x), v^2) &= \mathbb{E}_{\nu(Z)} \frac{|Y - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(X)|^2 + \text{Var}_{\theta|\mathbf{z}^N} f_\theta(X)}{2v^2} + \frac{\ln 2\pi v^2}{2} \\ &= \frac{\text{PR}^{(2)}(Y|X, \mathbf{z}^N) + \text{BER}^{(2)}(Y|X, \mathbf{z}^N)}{2v^2} + \frac{\ln 2\pi v^2}{2}. \end{aligned} \quad (19)$$

Eq. (19) implies that the standard VI tries to fit the mean of the predictive distribution to target variable y with the regularization term about the variance of the predictor. These terms corresponds

to $\text{PR}^{(2)}(Y|X, \mathbf{z}^N)$ and $\text{BER}^{(2)}(Y|X, \mathbf{z}^N)$. Note that there is a relation $\text{BER}^{\log}(Y|x, \mathbf{z}^N) \leq \text{Var}_{\theta|\mathbf{z}^N} f_{\theta}(x)/v^2$ for the Gaussian likelihood. This interpretation is consistent with Remark 1.

It has been numerically reported that the standard VI often underestimates EU. Alternative objective functions have been proposed to address this issue. For example, the entropic loss defined as $\text{Ent}_{\alpha}^l(y, x) := -\frac{1}{\alpha} \ln \mathbb{E}_{q(\theta|\mathbf{z}^N)} e^{-\alpha l(y, f_{\theta}(x))}$ for $\alpha > 0$, which is used in the α -divergence dropout (α -DO) (Li & Gal, 2017) and the second order PAC-Bayesian methods (2nd-PAC) (Masegosa, 2020; Futami et al., 2021), can capture EU better than the standard VI. Note that when $\alpha = 1$ and the log loss is used, the entropic risk corresponds to the log loss using the predictive distribution. For the Gaussian likelihood, we can upper-bound the entropic risk:

$$\mathbb{E}_{\nu(Z)} \text{Ent}_{\alpha=1}^{\log}(Y, X) \leq \mathbb{E}_{\nu(Z)} \frac{|Y - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_{\theta}(X)|^2}{v^2} + \frac{\text{Var}_{\theta|\mathbf{z}^N} f_{\theta}(X)}{v^2} - \text{BER}^{\log}(Y|X, \mathbf{z}^N) + \ln 2\pi v^2, \quad (20)$$

where we used Eq. (17). See Appendix D.12 for the derivation. Compared to Eq. (19), the entropic risk implicitly introduces a smaller regularization term for BER. This explains why α -DO and 2nd-PAC showed larger EU than the standard VI. We show a similar result for the entropic risk of the general log loss other than the Gaussian likelihood in Appendix D.12.

From these relations, balancing BER and PR appropriately leads to a solution that better evaluates the EU. Motivated by the decomposition in Eqs. (19) and (20), we directly control the prediction performance and the Bayesian excess risk for the Gaussian likelihood:

$$\text{rBER}(\lambda) = \frac{1}{N} \sum_{i=1}^N \frac{|y_i - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_{\theta}(x_i)|^2}{2v^2} + \lambda \frac{\text{Var}_{\theta|\mathbf{z}^N} f_{\theta}(x_i)}{2v^2} + \frac{\ln 2\pi v^2}{2} + \frac{1}{N} \text{KL}(q(\theta|\mathbf{z}^N)|p(\theta)), \quad (21)$$

where $0 < \lambda \leq 1$ is the coefficient of the BER regularizer. $\text{KL}(q(\theta|\mathbf{z}^N)|p(\theta))$ is a regularization term motivated by PAC-Bayesian theory. We select λ by cross-validation and it should be smaller than 1 since $\lambda = 1$ corresponds to the standard VI from Eq. (19) and the standard VI often underestimates the EU. We call Eq. (21) the regularized Bayesian Excess Risk VI (rBER) and show the PAC-Bayesian generalization guarantee for our rBER in Appendix F. In Sec. 5.2, we numerically evaluated this objective function.

rBER can also be seen as an extension of the standard VI. In the standard VI, the test loss is lower-bounded by the sum of the PR and BER with equal weights (Eq. (15)). rBER has the flexible weights between PR and BER. See Appendix F for a detailed comparison.

4 RELATION TO EXISTING WORK

The existing theoretical analysis of uncertainty focused on the calibration performance and clarified when a model over- and underestimates uncertainty (Nixon et al.; Bai et al., 2021; Naeini et al., 2015; Guo et al., 2017). We did not focus on analyzing the calibration property in this work because considering only the EU is not sufficient to deal with the calibration. We need to study both the AU and EU simultaneously. This is the limitation of our work. Other than calibration, the analysis of Gaussian processes (GP) has been gaining attention since GP’s posterior predictive distribution can be expressed analytically (Fiedler et al., 2021; Lederer et al., 2019). Some research focused on the distance or geometry between the test and training data points to derive EU (Liu et al., 2020; Tian et al., 2021). Other approaches connect the randomness of the posterior distribution to predictions by the delta method (Nilsen et al., 2022). Differently, the information-theoretic approach (Xu & Raginsky, 2020) focused on the loss function of the problem and defined the excess risk as the EU. Loss function-based analysis was proposed in the deterministic learning algorithm (Jain et al., 2021). Our theory, which can be regarded as an extension of the information-theoretic approach (Xu & Raginsky, 2020) to approximate Bayesian inference, derived the convergence properties of the variance and the entropy of the posterior predictive distributions.

Although the excess risk bound in Eq. (2) has been discussed by PAC-Bayesian theory (Alquier, 2021), its relation to the EU has not been investigated. The relationship between PAC-Bayesian theory and Bayesian inference has been investigated in terms of marginal likelihood (Germain et al., 2016; Rothfuss et al., 2021). Our work established new relationships that connect the uncertainty of the Bayesian predictive distribution and the PAC-Bayesian generalization bound. The information-theoretic approach (Xu & Raginsky, 2020) clarified that EU can be expressed by conditional mutual information. This relation was extended to meta-learning (Jose et al., 2021). However, the researchers

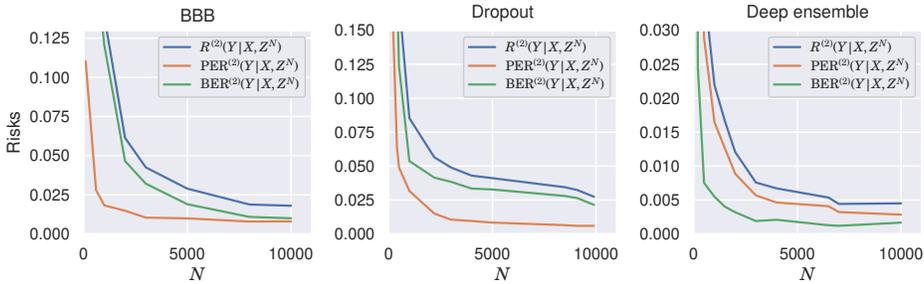


Figure 1: Result of toy data experiments: N represents number of training data points, and vertical line is value of each excess risk.

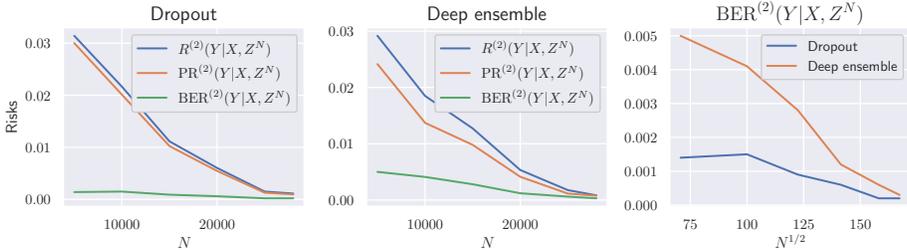


Figure 2: Real data experiments of depth estimation: Vertical line is the value of each risk.

assumed that correct models and exact posterior distributions are available. Our proposed analysis relaxes these assumptions.

5 NUMERICAL EXPERIMENTS

In this section, we numerically confirm the theoretical findings in Sec. 3 and our proposed rBER in Eq. (21). We show the detailed experimental settings and additional results in Appendix G.

5.1 NUMERICAL EVALUATION OF THEOREM 2

We numerically confirm the statement of Theorem 2. First, we consider toy data experiments where the true model is $y = 0.5x^3 + \epsilon$, $\epsilon \sim N(0, 1)$, $x \sim N(0, 1)$. We consider a Bayesian neural network (BNN) for $f_\theta(x)$ as a 4 layer neural network model with ReLU activation. We approximate the posterior distribution of the parameters of the neural network by Bayes by back-propagation (BBB)(Hernández-Lobato & Adams, 2015), dropout (Kendall & Gal, 2017), and deep ensemble (Lakshminarayanan et al., 2017). We evaluate $\text{PER}^{(2)}(Y|X, \mathbf{Z})$, $\text{BER}^{(2)}(Y|X, \mathbf{Z})$ ($:= \mathbb{E}_{\nu(X)} \text{Var}_{\theta|Z^N} f_\theta(X)$), and $R^{(2)}(Y|X, \mathbf{Z})$ (test error). The results are shown in Fig. 1. Our numerical results satisfy Eq. (13) in Theorem 2, that is, $\text{PER}^{(2)}(Y|X, \mathbf{Z})$ and $\text{BER}^{(2)}(Y|X, \mathbf{Z})$ are upper-bounded by $R^{(2)}(Y|X, \mathbf{Z})$ and converge to zero as the number of samples increases. We calculated the Spearman Rank Correlation (SRC) among $\text{PER}^{(2)}(Y|X, \mathbf{Z})$, $\text{BER}^{(2)}(Y|X, \mathbf{Z})$, and $R^{(2)}(Y|X, \mathbf{Z})$ and showed at least 0.97 suggesting high correlation relation between them.

Next, we confirm Theorem 2 using a real-world dataset. Following the setting of existing work (Amini et al., 2020), we trained a U-Net style network (Ronneberger et al., 2015) with the data of the NYU Depth v2 dataset (Silberman et al., 2012), which consists of RGB-to-depth. We applied dropout and deep ensemble methods. Since we cannot evaluate $\text{PER}^{(2)}(Y|X, \mathbf{Z})$, we instead evaluated Eq. (15), which only requires $\text{PR}^{(2)}(Y|X, \mathbf{Z})$, $\text{BER}^{(2)}(Y|X, \mathbf{Z})$, and $R^{(2)}(Y|X, \mathbf{Z})$. The result is shown in Fig. 2. We found that $\text{PR}^{(2)}(Y|X, \mathbf{Z})$, $\text{BER}^{(2)}(Y|X, \mathbf{Z})$ are upper-bounded by $R^{(2)}(Y|X, \mathbf{Z})$ for real dataset experiments. We calculated the SRC among $\text{PR}^{(2)}(Y|X, \mathbf{Z})$, $\text{BER}^{(2)}(Y|X, \mathbf{Z})$, and

Table 1: Benchmark results on test RMSE, PICP, and MPIW.

Dataset	Avg. Test RMSE				Avg. Test PICP and MPIW in parenthesis			
	f-SVGD	VAR	rBER(0)	rBER(0.05)	f-SVGD	VAR	rBER(0)	rBER(0.05)
Concrete	4.33±0.8	4.30±0.7	4.47±0.6	4.48±0.7	0.82±0.03 (0.13±0.00)	0.87±0.04 (0.16±0.01)	0.99±0.02 (0.50±0.04)	0.95±0.02 (0.25±0.02)
Boston	2.54±0.50	2.53±0.50	2.53±0.50	2.53±0.51	0.63±0.07 (0.10±0.02)	0.76±0.05 (0.14±0.01)	0.97±0.01 (0.33±0.04)	0.92±0.04 (0.22±0.02)
Wine	0.61±0.04	0.61±0.04	0.64±0.04	0.63±0.02	0.79±0.03 (0.32±0.05)	0.85±0.02 (0.39±0.06)	0.99±0.00 (1.61±0.00)	0.95±0.03 (0.32±0.15)
Power	3.78±0.14	3.75±0.13	3.66±0.15	3.69±0.12	0.43±0.01 (0.07±0.00)	0.82±0.01 (0.15±0.00)	0.99±0.01 (0.81±0.01)	0.96±0.01 (0.37±0.01)
Yacht	0.64±0.28	0.60±0.28	0.75±0.41	0.78±0.48	0.92±0.04 (0.02±0.01)	0.93±0.04 (0.04±0.01)	0.96±0.03 (0.10±0.01)	0.94±0.04 (0.08±0.01)
Protein	3.98±0.54	3.92±0.05	3.83±0.10	3.85±0.05	0.53±0.01 (0.24±0.01)	0.83±0.00 (0.58±0.01)	1.0 ±0.00 (5.04±0.01)	0.96±0.01 (0.86±0.00)

Table 2: Cumulative regret relative to that of the uniform sampling.

Dataset	MAP	PAC _E ²	f-SVGD	VAR	rBER(0)	rBER(0.01)	rBER(0.05)
Mushroom	0.129±0.098	0.037±0.012	0.043±0.009	0.029±0.010	0.075±0.005	0.024±0.009	0.021±0.004
Financial	0.791±0.219	0.189±0.025	0.154±0.017	0.155±0.024	0.351±0.030	0.075±0.024	0.075±0.031
Statlog	0.675 ±0.287	0.032±0.003	0.010±0.000	0.006±0.000	0.145±0.223	0.005±0.001	0.005±0.000
CoverType	0.610±0.051	0.396±0.006	0.372±0.007	0.291±0.004	0.610±0.051	0.351±0.003	0.290±0.002

$R^{(2)}(Y|X, \mathbf{Z})$ and showed at least 0.98, suggesting a high correlation relation between them. We also evaluated the convergence behaviors of BER and show the result on the right in Fig. 2. BER converges with $\mathcal{O}(1/N^{1/2})$, which is consistent with Eq. (13).

5.2 REAL DATA EXPERIMENTS OF REGULARIZED BAYESIAN EXCESS RISK VI

We numerically compared the prediction and EU evaluation performances of our proposed method shown in Eq. (21) in regression and contextual bandit tasks. Motivated by the success of the entropic risk in particle VI (PVI) (Masegosa, 2020; Futami et al., 2021), which approximates the posterior distribution by the ensemble of models, we also applied our rBER to the PVI setting. Thus, the posterior distribution is expressed as $q(\theta) := \frac{1}{N} \sum_{i=1}^M \delta_{\theta_i}(\theta)$, where $\delta_{\theta_i}(\theta)$ is the Dirac distribution that has a mass at θ_i . See Appendix G for details about PVI. We refer to rBER(0) when $\lambda = 0$ in Eq. (21). We compared our method with the existing PVI methods, f-SVGD (Wang et al., 2019), PAC_E² (Masegosa, 2020), and VAR (Futami et al., 2021).

We used the UCI dataset (Dheeru & Karra Taniskidou, 2017) for regression tasks. The model is a single-layer network with ReLU activation, and we used 20 ensembles. The results of 20 repetitions are shown in Table 1. We evaluated the fitting performance by RMSE and the uncertainty estimation performance by the prediction interval coverage probability (PICP), which shows the number of test observations inside the estimated prediction interval where the interval was set to 0.95. PICP is best when it is close to 0.95. We evaluated the mean prediction interval width (MPIW), which shows an average width of a prediction interval. A smaller MPIW is a better uncertainty estimate when PICP is near the best. Due to space limitations, the results of PAC_E² and the other λ s and the negative log-likelihood are shown in Appendix G. We found the existing PVIs show small PICP and MPIW, indicating that the existing methods underestimate the uncertainty. rBER(0) shows a large PICP and MPIW since the Bayesian excess risk is not regularized. rBER(0.05) shows a moderate MPIW with a better PICP and almost identical prediction performance in RMSE. Thus, rBER successfully controlled the prediction and the uncertainty evaluation performances.

Next we evaluated the rBER using contextual bandit problems (Riquelme et al., 2018). We need to balance the trade-off between exploitation and exploration to achieve small cumulative regret. For that purpose, our algorithms must appropriately control the prediction and uncertainty evaluation performance. We used the Thompson sampling algorithm with BNN and two hidden layers. We used 20 ensembles for approximating the posterior distribution. The results of 10 repetitions are shown in Table 2. Our approach outperformed other methods, which means our proposed method showed better prediction and uncertainty control than the existing methods.

6 CONCLUSION

We theoretically and numerically analyzed the epistemic uncertainty of approximate inference. We clarified the novel relations among excess risk, epistemic uncertainty, and generalization error. We then showed the convergence rate of the widely used uncertainty measures for the first time. Motivated by theoretical analysis, we proposed a novel variational inference (VI) and applied it to the particle VI. In future work, it would be interesting to explore the relation between BER and evidential learning.

REFERENCES

- Pierre Alquier. User-friendly introduction to pac-bayes bounds. *arXiv preprint arXiv:2110.11216*, 2021.
- Pierre Alquier, James Ridgway, and Nicolas Chopin. On the properties of variational approximations of gibbs posteriors. *The Journal of Machine Learning Research*, 17(1):8374–8414, 2016.
- Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. *Advances in Neural Information Processing Systems*, 33:14927–14937, 2020.
- Yu Bai, Song Mei, Huan Wang, and Caiming Xiong. Don’t just blame over-parametrization for over-confidence: Theoretical analysis of calibration in binary classification. In *International Conference on Machine Learning*, pp. 566–576. PMLR, 2021.
- Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, et al. Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 401–413, 2021.
- Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Lawrence D Brown and Roger Purves. Measurable selections of extrema. *The annals of statistics*, pp. 902–912, 1973.
- B.S. Clarke and A.R. Barron. Information-theoretic asymptotics of bayes methods. *IEEE Transactions on Information Theory*, 36(3):453–471, 1990. doi: 10.1109/18.54897.
- Stefan Depeweg, Jose-Miguel Hernandez-Lobato, Finale Doshi-Velez, and Steffen Udluft. Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning. In *International Conference on Machine Learning*, pp. 1184–1193. PMLR, 2018.
- Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Christian Fiedler, Carsten W Scherer, and Sebastian Trimpe. Practical and rigorous uncertainty bounds for gaussian process regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 7439–7447, 2021.
- Adam Foster, Martin Jankowiak, Elias Bingham, Paul Horsfall, Yee Whye Teh, Thomas Rainforth, and Noah Goodman. Variational bayesian optimal experimental design. *Advances in Neural Information Processing Systems*, 32, 2019.
- Futoshi Futami, Tomoharu Iwata, Issei Sato, Masashi Sugiyama, et al. Loss function based second-order jensen inequality and its application to particle variational inference. *Advances in Neural Information Processing Systems*, 34, 2021.
- Jakob Gawlikowski, Cedrique Rovile Njiteutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, Muhammad Shahzad, Wen Yang, Richard Bamler, and Xiao Xiang Zhu. A survey of uncertainty in deep neural networks, 2022.
- Pascal Germain, Francis Bach, Alexandre Lacoste, and Simon Lacoste-Julien. Pac-bayesian theory meets bayesian inference. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 1884–1892, 2016.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning*, pp. 1321–1330. PMLR, 2017.
- Hassan Hafez-Kolahi, Behrad Moniri, Shohreh Kasaei, and Mahdieh Soleymani Baghshah. Rate-distortion analysis of minimum excess risk in bayesian learning. In *International Conference on Machine Learning*, pp. 3998–4007. PMLR, 2021.

- David Haussler and Manfred Opper. Mutual information, metric entropy and cumulative relative entropy risk. *The Annals of Statistics*, 25(6):2451–2492, 1997.
- Rianne Heide, Alisa Kirichenko, Peter Grunwald, and Nishant Mehta. Safe-bayesian generalized linear regression. In *International Conference on Artificial Intelligence and Statistics*, pp. 2623–2633. PMLR, 2020.
- José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International conference on machine learning*, pp. 1861–1869. PMLR, 2015.
- José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. *Advances in neural information processing systems*, 27, 2014.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3):457–506, 2021.
- Moksh Jain, Salem Lahlou, Hadi Nekoei, Victor Butoi, Paul Bertin, Jarrid Rector-Brooks, Maksym Korablyov, and Yoshua Bengio. Deup: Direct epistemic uncertainty prediction. *arXiv preprint arXiv:2102.08501*, 2021.
- David Janz, Jiri Hron, Przemysław Mazur, Katja Hofmann, José Miguel Hernández-Lobato, and Sebastian Tschiatschek. Successor uncertainties: exploration and uncertainty in temporal difference learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Sharu Theresa Jose, Sangwoo Park, and Osvaldo Simeone. Information-theoretic analysis of epistemic uncertainty in bayesian meta-learning. *arXiv preprint arXiv:2106.00252*, 2021.
- Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Armin Lederer, Jonas Umlauft, and Sandra Hirche. Posterior variance analysis of gaussian processes with application to average learning curves, 2019.
- Yingzhen Li and Yarin Gal. Dropout inference in bayesian neural networks with alpha-divergences. In *International conference on machine learning*, pp. 2052–2061. PMLR, 2017.
- Jeremiah Liu, Zi Lin, Shreyas Padhy, Dustin Tran, Tania Bedrax Weiss, and Balaji Lakshminarayanan. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. *Advances in Neural Information Processing Systems*, 33:7498–7512, 2020.
- Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances In Neural Information Processing Systems*, pp. 2378–2386, 2016.
- Andres Masegosa. Learning under model misspecification: Applications to variational and ensemble methods. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 5479–5491. Curran Associates, Inc., 2020.
- Andres Masegosa, Stephan Lorenzen, Christian Igel, and Yevgeny Seldin. Second order pac-bayesian bounds for the weighted majority vote. *Advances in Neural Information Processing Systems*, 33, 2020.
- Mahdi Pakdaman Naeni, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

- Geir K Nilsen, Antonella Z Munthe-Kaas, Hans J Skaug, and Morten Brun. Epistemic uncertainty quantification in deep learning classification by the delta method. *Neural Networks*, 145:164–176, 2022.
- Jeremy Nixon, Michael W Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. Measuring calibration in deep learning.
- Yuki Ohnishi and Jean Honorio. Novel change of measure inequalities with applications to pac-bayesian bounds and monte carlo estimation. In *International Conference on Artificial Intelligence and Statistics*, pp. 1711–1719. PMLR, 2021.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. In *International Conference on Learning Representations*, 2018.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.
- Jonas Rothfuss, Vincent Fortuin, Martin Josifoski, and Andreas Krause. Pacoh: Bayes-optimal meta-learning with pac-guarantees. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9116–9126. PMLR, 18–24 Jul 2021.
- Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European conference on computer vision*, pp. 746–760. Springer, 2012.
- Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- Junjiao Tian, Dylan Yung, Yen-Chang Hsu, and Zsolt Kira. A geometric perspective towards neural calibration via sensitivity decomposition. *Advances in Neural Information Processing Systems*, 34, 2021.
- Ziyu Wang, Tongzheng Ren, Jun Zhu, and Bo Zhang. Function space particle optimization for bayesian neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=BkgtDsCcKQ>.
- Sumio Watanabe. *Algebraic geometry and statistical learning theory*. Number 25. Cambridge university press, 2009.
- Sumio Watanabe. *Mathematical theory of Bayesian statistics*. Chapman and Hall/CRC, 2018.
- Aolin Xu. Continuity of generalized entropy and statistical learning. *arXiv preprint arXiv:2012.15829*, 2020.
- Aolin Xu and Maxim Raginsky. Minimum excess risk in bayesian learning. *arXiv preprint arXiv:2012.14868*, 2020.
- Nanyang Ye and Zhanxing Zhu. Bayesian adversarial learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

A NOTATION

Distributions

$\nu(z)$	A data generating distribution
$p(y x, \theta)$	A model
$p(\theta)$	A prior distribution
$q(\theta \mathbf{z}^N)$	A posterior distribution
$p^q(y x, \mathbf{z}^N)$	The predictive distribution obtained by the expectation over $q(\theta \mathbf{z}^N)$
$p^q(\theta, \mathbf{z}^N, z)$	The approximate joint distribution defined as $\nu(\mathbf{z}^N)q(\theta \mathbf{z}^N)\nu(x)p(y x, \theta)$
$p_B(\mathbf{z}^N, z, \theta)$	The joint model used in Xu & Raginsky (2020) defined as $p(\theta) \prod_{n=1}^N p(y_n x_n, \theta)\nu(x_n)p(y x, \theta)\nu(x)$.

Risk functions

$R^l(Y X, \mathbf{Z}^N)$	A test error defined as $\mathbb{E}_{\nu(\mathbf{Z}^N)}\mathbb{E}_{q(\theta \mathbf{Z}^N)}\mathbb{E}_{\nu(Z)}l(Y, f_\theta(X))$
$\mathbb{E}_{\nu(\mathbf{Z}^N)}r^l(\mathbf{Z}^N)$	A training error defined as $\mathbb{E}_{\nu(\mathbf{Z}^N)}\mathbb{E}_{q(\theta \mathbf{Z}^N)}\sum_{n=1}^N l(Y_n, f_\theta(X_n))/N$
$\text{ER}^l(Y X, \mathbf{Z}^N, \theta^*)$	The excess risk defined as $R^l(Y X, \mathbf{Z}^N) - R^l(Y X, \theta^*)$.
$\text{PR}^l(Y X, \mathbf{Z}^N)$	A prediction risk defined as $\mathbb{E}_{\nu(\mathbf{Z}^N)}\mathbb{E}_{\nu(Z)}l(Y, \mathbb{E}_{q(\theta \mathbf{Z}^N)}f_\theta(X))$
$\text{PER}^l(Y X, \mathbf{Z}^N)$	The prediction excess risk defined as $\text{PR}^l(Y X, \mathbf{Z}^N) - \inf_{\phi: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Z)}[l(Y, \phi(X))]$
$\text{BPR}^l(Y X, \mathbf{Z}^N)$	The Bayesian prediction risk defined as $\mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)}l(Y, \mathbb{E}_{q(\theta \mathbf{Z}^N)}f_{\theta'}(X))$
$\text{BER}^l(Y X, \mathbf{Z}^N)$	The defined Bayesian excess risk as $\text{BPR}^l(Y X, \mathbf{Z}^N) - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)}l(Y, \phi(\theta, X))$.
$\text{MER}^l(Y X, \mathbf{Z}^N)$	A minimum excess risk used in Xu & Raginsky (2020)

B SUMMARY OF SETTINGS

Here we summarize the concepts and definitions of joint distributions and risks used in this work.

B.1 BAYESIAN LEARNING (SEC 2.3) [USED IN (XU & RAGINSKY, 2020) AND (JOSE ET AL., 2021)]

- Joint distribution (All the data is conditionally i.i.d.):

$$p_B(\mathbf{Z}^N, Z, \theta) := p(\theta) \prod_{n=1}^N p(Y_n|X_n, \theta)\nu(X_n)p(Y|X, \theta)\nu(X).$$

- Minimum excess risk:

$$\begin{aligned} \text{MER}^l(Y|X, \mathbf{Z}^N) &:= R^l(Y|X, \mathbf{Z}^N) - R^l(Y|X, \theta) \\ R^l(Y|X, \mathbf{Z}^N) &:= \inf_{\psi: \mathbf{Z}^N \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p_B(\mathbf{Z}^N, Z, \theta)}[l(Y, \psi(X, \mathbf{Z}^N))], \\ R^l(Y|X, \theta) &:= \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p_B(\mathbf{Z}^N, Z, \theta)}[l(Y, \phi(\theta, X))]. \end{aligned}$$

B.2 THE SETTING USED IN THE PAC-BAYESIAN THEORY (SEC 2.1)

- Joint distribution of data and parameter θ (All the data is i.i.d.):

$$\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(Z).$$

- Prediction excess risk:

$$\begin{aligned} \text{PER}^l(Y|X, \mathbf{Z}^N) &:= \text{PR}^l(Y|X, \mathbf{Z}^N) - \inf_{\phi: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Z)}[l(Y, \phi(X))], \\ \text{PR}^l(Y|X, \mathbf{Z}^N) &:= \mathbb{E}_{\nu(\mathbf{Z}^N)} \mathbb{E}_{\nu(Z)} l(Y, \mathbb{E}_{q(\theta|\mathbf{Z}^N)} f_{\theta}(X)). \end{aligned}$$

B.3 OUR SETTING DEFINED IN SEC 3.1)

- Joint distribution (The training data is i.i.d. The test data follows the model):

$$p^q(\theta, \mathbf{Z}^N, Z) := \nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X)p(Y|X, \theta)$$

- Bayesian excess risk:

$$\begin{aligned} \text{BER}^l(Y|X, \mathbf{Z}^N) &:= \text{BPR}^l(Y|X, \mathbf{Z}^N) - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)} l(Y, \phi(\theta, X)), \\ \text{BPR}^l(Y|X, \mathbf{Z}^N) &:= \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)} l(Y, \mathbb{E}_{q(\theta'|\mathbf{Z}^N)} f_{\theta'}(X)). \end{aligned}$$

C FURTHER PRELIMINARIES

C.1 ADDITIONAL FACTS ABOUT BAYESIAN LEARNING

Here we introduce the preliminary results (Xu & Raginsky, 2020) about the MER in a Bayesian setting. Besides the log loss, we can upper bound MER by conditional mutual information. First, we can upper bound MER by the plug-in decision rule. Consider an optimal decision rule $\Psi^* : \mathcal{Z} \times \Theta \rightarrow \mathcal{A}$ and this satisfies $R^l(Y|X, \theta) = \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\mathbf{Z}^N, Z, \theta)} [l(Y, \phi(\theta, X))] = \mathbb{E} l(Y, \Psi^*(X, \theta))$. Then we express θ' is drawn from a posterior distribution $p(\theta|\mathbf{Z}^N)$. Then we have

$$\text{MER}_l(Y|X, \mathbf{Z}^N) \leq \mathbb{E} l(Y, \Psi^*(X, \theta')) - l(Y, \Psi^*(X, \theta)), \quad (22)$$

where $\Psi^*(X, \theta')$ is a plug-in decision rule, first we draw θ' from posterior distribution and substitute it to Ψ^* . Then if the moment generating function of $l(Y, \Psi^*(X, \theta'))$ under $P(Y, \theta'|X, \mathbf{Z}^N)$ satisfies regularity conditions, we can upper bound MER. For example, a loss function satisfies σ^2 -subGaussian conditioned on $(X, \mathbf{z}^N) = (x, \mathbf{z}^N)$ for all (x, \mathbf{z}^N) , then

$$\text{MER}_l(Y|X, \mathbf{Z}^N) = \sqrt{2\sigma^2 I(\theta; Y|X, \mathbf{Z}^N)}. \quad (23)$$

Thus, we can treat zero-one loss and some squared loss.

Thus, CMI plays a central role in the Bayesian excess risk analysis. Then existing work (Xu & Raginsky, 2020) shows that

$$I(\theta; Y|X, \mathbf{Z}^N) \leq \frac{1}{N} I(\theta; \mathbf{Z}^N). \quad (24)$$

This is because the mutual information is upper-bounded by $\mathcal{O}(\ln N)$ for many practical settings. Thus, the excess risk converges to 0 as $N \rightarrow \infty$.

C.2 PRELIMINARIES OF THE PAC-BAYESIAN THEORY

We briefly introduce the PAC-Bayesian theory. The typical PAC-Bayesian bound provides us the high-probability guarantee about the gap between the test error $\tilde{R}(\theta) := \mathbb{E}_{\nu(Z)} l(Y, f_{\theta}(X))$ and $\tilde{r}(\theta) := \frac{1}{N} \sum_{n=1}^N l(Y_n, f_{\theta}(X_n))$ (Here we do not take the expectation over \mathbf{Z}^N);

Theorem 4. (Alquier et al., 2016) *Given a data generating distribution ν , for any prior distribution $p(\theta)$ over Θ independent of \mathbf{Z}^N and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over the choice of training data $\nu(\mathbf{Z}^N)$, for all probability distributions $q(\theta|\mathbf{z}^N)$ over Θ , we have*

$$\mathbb{E}_q \tilde{R}(\theta) \leq \mathbb{E}_q \tilde{r}(\theta) + \frac{\text{KL}(q|p) + \ln \xi^{-1} + \Omega_{p,\nu}(c, N)}{cN}, \quad (25)$$

where $\Omega_{p,\nu}(c, N) := \ln \mathbb{E}_{p(\theta)} \mathbb{E}_{\nu(\mathbf{Z}^N)} \exp[cN(\tilde{R}(\theta) - \tilde{r}(\theta))]$.

This constant $\Psi_{p,\nu}$ depends on the property of the loss function and the data generating distribution and prior. For example, when $l(y, f_\theta(x)) - \mathbb{E}_\nu l(Y, f_\theta(X))$ satisfies the σ^2 sub-Gaussian property, and by setting $c = 1/\sqrt{N}$, we have

$$\mathbb{E}_q \tilde{R}(\theta) \leq \mathbb{E}_q \tilde{r}(\theta) + \frac{\text{KL}(q|p) + \ln \xi^{-1} + \sigma^2/2}{\sqrt{N}}. \quad (26)$$

On the other hand, we introduced the bound in expectation in the main paper (Alquier, 2021). For example, under the similar setting as Eq. (26), when we assume the σ^2 sub-Gaussian property and $c = \lambda$, we have

$$R^l(Y|X, \mathbf{Z}^N) \leq r(\theta) + \frac{\text{KL}(q(\theta|\mathbf{Z}^N)|p(\theta))}{\lambda} + \frac{\lambda\sigma^2}{2N}, \quad (27)$$

see Alquier (2021) for the proof and other settings.

Next, we introduce the PAC-Bayesian bound Eq. (2). When $r(\theta)$ satisfies the L -lipschitz property and setting the prior as $N(\theta|0, \beta^2 I_d)$ and $\lambda = 1/\sqrt{N}$, we have

$$R^l(Y|X, \mathbf{Z}^N) - R^l(Y|X, \theta^*) \leq L\beta\sqrt{\frac{d}{N}} + \frac{\sigma^2}{2\sqrt{N}} + \frac{\frac{\|\theta^*\|^2}{2\beta^2} + \frac{d}{2}\log N}{\sqrt{N}}. \quad (28)$$

See (Alquier, 2021) for the proof and other settings.

In the main paper, we considered the squared loss. For the squared loss, above Lipschitz bound cannot be used. Here we introduce different PAC-Bayesian bound. In stead of Lipschitz property, we assume that for any $\theta, \theta' \in \Theta$, there exists a measurable function $M(x)$ such that

$$f_\theta(x) - f_{\theta'}(x) \leq M(x)\|\theta - \theta'\|^2, \quad (29)$$

and assume $\mathbb{E}_{\nu(X)} M(X) < L < \infty$. From the almost identical derivation of Example 2.2 in Alquier (2021), we have

$$R^l(Y|X, \mathbf{Z}^N) - R^l(Y|X, \theta^*) \leq L\beta^2 \frac{d}{N} + \frac{\sigma^2}{2\sqrt{N}} + \frac{\frac{\|\theta^*\|^2}{2\beta^2} + \frac{d}{2}\log N}{\sqrt{N}}. \quad (30)$$

D PROOFS OF THEOREMS IN SECTION 3

Here we present the proofs of Section 3.

D.1 CONDITIONAL EXPECTATION OF EXCESS RISKS

In this section, we define the conditional version of the excess risks. Note that PER and BER was defined as

$$\begin{aligned} \text{PER}^l(Y|X, \mathbf{Z}^N) &:= \text{PR}^l(Y|X, \mathbf{Z}^N) - \inf_{\tilde{\phi}: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Z)} [l(Y, \tilde{\phi}(X))], \\ \text{BER}^l(Y|X, \mathbf{Z}^N) &:= \text{BPR}^l(Y|X, \mathbf{Z}^N) - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)} l(Y, \phi(\theta, X)). \end{aligned} \quad (31)$$

We define the conditional excess risk as

$$\text{PER}^l(Y|x, \mathbf{z}^N) := \text{PR}^l(Y|x, \mathbf{z}^N) - \inf_{\tilde{\phi}: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Y|x)} [l(Y, \tilde{\phi}(x))], \quad (32)$$

$$\text{BER}^l(Y|x, \mathbf{z}^N) := \text{BPR}^l(Y|x, \mathbf{z}^N) - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{q(\theta|\mathbf{z}^N)p(Y|x,\theta)} l(Y, \phi(\theta, x)). \quad (33)$$

It has been proved in Lemma 3.4 in Steinwart & Christmann (2008) that if the action space is \mathbb{R} following relation holds

$$\mathbb{E}_{\nu(\mathbf{Z}^N = \mathbf{z}^N)\nu(X=x)} \text{PER}^l(Y|x, \mathbf{z}^N) = \text{PER}^l(Y|X, \mathbf{Z}^N), \quad (34)$$

$$\mathbb{E}_{\nu(\mathbf{Z}^N = \mathbf{z}^N)\nu(X=x)} \text{BER}^l(Y|x, \mathbf{z}^N) = \text{BER}^l(Y|X, \mathbf{Z}^N). \quad (35)$$

Moreover, for the log loss, from Theorem 3 in Brown & Purves (1973), above relation holds. Thus, we can naturally connect the conditional and unconditional definitions of PER and BER for the squared loss and log-loss.

In the following, we explicitly calculate how relations Eqs. (34) and (35) holds. The first terms in (32) and (33) can easily be expressed as

$$\begin{aligned} \text{PR}^l(Y|X, \mathbf{Z}^N) &= \mathbb{E}_{\nu(\mathbf{Z}^N)} \mathbb{E}_{\nu(Z)} l(Y, \mathbb{E}_{q(\theta|\mathbf{Z}^N)} f_{\theta}(X)) \\ &= \mathbb{E}_{\nu(\mathbf{Z}^N=\mathbf{z}^N)\nu(X=x)} \mathbb{E}_{\nu(Y|X=x)} l(Y, \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_{\theta}(x)) \\ &= \mathbb{E}_{\nu(\mathbf{Z}^N=\mathbf{z}^N)\nu(X=x)} \text{PR}^l(Y|x, \mathbf{z}^N), \end{aligned} \quad (36)$$

and

$$\begin{aligned} \text{BPR}^l(Y|X, \mathbf{Z}^N) &= \mathbb{E}_{\nu(\mathbf{Z}^N)} \mathbb{E}_{\nu(X)q(\theta|\mathbf{Z}^N)p(Y|X,\theta)} l(Y, \mathbb{E}_{q(\theta|\mathbf{Z}^N)} f_{\theta}(X)) \\ &= \mathbb{E}_{\nu(\mathbf{Z}^N=\mathbf{z}^N)\nu(X=x)} \mathbb{E}_{q(\theta|\mathbf{z}^N)p(Y|X=x,\theta)} l(Y, \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_{\theta}(x)) \\ &= \mathbb{E}_{\nu(\mathbf{Z}^N=\mathbf{z}^N)\nu(X=x)} \text{BPR}^l(Y|x, \mathbf{z}^N). \end{aligned} \quad (37)$$

Next we calculate the Bayes risks in the second terms. First, for the squared loss, we have

$$\begin{aligned} \mathbb{E}_{\nu(Z)} [l(Y, \tilde{\phi}(X))] &= \mathbb{E}_{\nu(Z)} (Y - \tilde{\phi}(X))^2 \\ &= \mathbb{E}_{\nu(Z)} (Y - \mathbb{E}_{\nu(Y'|X)} Y'|X)^2 + \mathbb{E}_{\nu(X)} (\mathbb{E}_{\nu(Y'|X)} Y'|X - \tilde{\phi}(X))^2, \end{aligned} \quad (38)$$

where $\mathbb{E}_{\nu(Y'|X)} Y'|X$ is the conditional expectation. Thus, infimum is achieved by setting $\mathbb{E}_{\nu(Y'|X)} Y'|X = \tilde{\phi}(X)$. Thus, we have

$$\begin{aligned} \inf_{\tilde{\phi}: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Z)} [l(Y, \tilde{\phi}(X))] &= \mathbb{E}_{\nu(Z)} (Y - \mathbb{E}_{\nu(Y'|X)} Y'|X)^2 \\ &= \mathbb{E}_{\nu(X=x)} \mathbb{E}_{\nu(Y|X=x)} (Y - \mathbb{E}_{\nu(Y'|x)} Y'|x)^2 \\ &= \mathbb{E}_{\nu(X=x)} \inf_{\tilde{\phi}: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Y|x)} [l(Y, \tilde{\phi}(x))]. \end{aligned} \quad (39)$$

We can show the same statement as

$$\inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)} l(Y, \phi(\theta, X)) = \mathbb{E}_{\nu(\mathbf{Z}^N=\mathbf{z}^N)\nu(X=x)} \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{q(\theta|\mathbf{z}^N)p(Y|x,\theta)} l(Y, \phi(\theta, x)). \quad (40)$$

Combined these relations, we get Eqs. (34).

Next, we discuss the conditional Bayes risk for the log-loss. We can proceed the calculation in the same way as the squared loss. Then we have

$$\inf_{\tilde{\phi}: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Z)} [l(Y, \tilde{\phi}(X))] = -\mathbb{E}_{\nu(Y|X)\nu(X)} \log \nu(Y|X) = \mathbb{E}_{\nu(X=x)} \inf_{\tilde{\phi}: \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{\nu(Y|x)} [l(Y, \tilde{\phi}(x))]. \quad (41)$$

The Bayes risk in BPR has the similar relation. Thus, it is clear that the relations shown in Eqs. (34) and (35) hold.

D.2 PROOF OF $\text{BER}^l(Y|X, \mathbf{Z}^N) \geq 0$

We remark that for the squared loss and log-loss, BPR can be written as

$$\mathbb{E}_{p^f(\theta, \mathbf{Z}^N, Z)} l(Y, \mathbb{E}_{q(\theta'|\mathbf{Z}^N)} f_{\theta'}(X)) = \inf_{\psi: \mathcal{Z}^N \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\mathbf{Z}^N, Z, \theta)} [l(Y, \psi(X, \mathbf{Z}^N))]. \quad (42)$$

This expression is similar to the definition of the first term in MER. Then, applying the same technique in Lemma 1 (Xu & Raginsky, 2020), we can show that $\text{BPR}^l(Y|X, \mathbf{Z}^N)$ satisfies the data processing inequality. Then given a Markov chain, for example, $(X, \mathbf{Z}^N) - (X, Z^{N+1}) - Y$, then we have $\text{BPR}^l(Y|X, \mathbf{Z}^N) \geq \text{BPR}^l(Y|X, Z^{N+1})$. Consider a Markov chain $(X, \mathbf{Z}^N) - (X, \theta) - Y$. Then by the data processing inequality, we have $\text{BPR}^l(Y|X, \mathbf{Z}^N) \geq \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{p^q(\theta, \mathbf{Z}^N, Z)} l(Y, \phi(\theta, X))$ since the Bayes error uses the parameter of $p^q(\mathbf{z}^N, \theta, z)$ directly. This concludes the proof.

D.3 PROOF OF LEMMA 1

By definition,

$$\begin{aligned}
\mathbb{E}_{q(\theta|\mathbf{z}^N)}\|y - f_\theta(x)\|^2 &= y^2 - 2y\mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x) + \mathbb{E}_{q(\theta|\mathbf{z}^N)}[f_\theta(x)^2] \\
&= (y - \mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x))^2 + \mathbb{E}_{q(\theta|\mathbf{z}^N)}[f_\theta(x)^2] - [\mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x)]^2 \\
&= (y - \mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x))^2 + \mathbb{E}_{q(\theta|\mathbf{z}^N)}[f_\theta(x) - \mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x)]^2 \\
&= (y - \mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x))^2 + \text{Var}f_\theta(x).
\end{aligned} \tag{43}$$

This concludes the proof.

D.4 PROOF OF THEOREM 1

Here we consider the conditional quantities of them. The formal definitions of conditional fundamental limit of learning and total risk are given in Appendix D.1.

For the log loss, we use the property of the entropy. For any probability distributions p and q , the entropy satisfies $H[p] := -\mathbb{E}_p \ln p = -\inf_q \mathbb{E}_p \ln q$. Then, by definition, it is clear that

$$\text{BPR}^{\log}(Y|x, \mathbf{z}^N) = -\mathbb{E}_{q(\theta|\mathbf{z}^N)}\mathbb{E}_{p(Y|x, \theta)} \ln \mathbb{E}_{q(\theta'|\mathbf{z}^N)}p(Y|x, \theta') = H[p^q(Y|x, \mathbf{z}^N)], \tag{44}$$

where $p^q(Y|x, \mathbf{z}^N) = \mathbb{E}_{q(\theta|\mathbf{z}^N)}p(Y|x, \theta)$ is the conditional predictive distribution. Also, by definition, the Bayes risk for the log loss is given as

$$\begin{aligned}
\inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{q(\theta|\mathbf{z}^N)p(Y|x, \theta)} l(Y, \phi(\theta, x)) &= -\inf_{p'} \mathbb{E}_{q(\theta|\mathbf{z}^N)p(Y|x, \theta)} [\ln p'(Y|x, \theta)] \\
&= -\mathbb{E}_{q(\theta|\mathbf{z}^N)}\mathbb{E}_{p(Y|x, \theta)} [\ln p(Y|x, \theta)] \\
&= E_{q(\theta|\mathbf{z}^N)}H[p(Y|x, \theta)].
\end{aligned} \tag{45}$$

Thus,

$$\text{BER}^{\log}(Y|x, \mathbf{z}^N) = \text{BPR}^{\log}(Y|x, \mathbf{z}^N) - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{q(\theta|\mathbf{z}^N)p(Y|x, \theta)} l(Y, \phi(\theta, x)) = I_\nu(\theta; Y|x, \mathbf{z}^N). \tag{46}$$

Next, for the squared loss, recall that for any random variable Y with distribution p , we have

$$\inf_a \mathbb{E}_p|Y - a|^2 = \mathbb{E}_p|Y - \mathbb{E}_{p(Y)}Y'|^2 = \text{Var}(Y). \tag{47}$$

Using this relation, we have

$$\begin{aligned}
&\text{BER}^{(2)}(Y|x, \mathbf{z}^N) \\
&= \text{BPR}^{(2)}(Y|x, \mathbf{z}^N) - \inf_{\phi: \Theta \times \mathcal{X} \rightarrow \mathcal{A}} \mathbb{E}_{q(\theta|\mathbf{z}^N)p(Y|x, \theta)} l(Y, \phi(\theta, x)) \\
&= \mathbb{E}_{q(\theta|\mathbf{z}^N)}\mathbb{E}_{p(Y|x, \theta)}\|Y - \mathbb{E}_{q(\theta'|\mathbf{z}^N)}f_{\theta'}(x)\|^2 - \mathbb{E}_{q(\theta|\mathbf{z}^N)}\mathbb{E}_{p(Y|x, \theta)}\|Y - f_\theta(x)\|^2 \\
&= -2\mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x)\mathbb{E}_{q(\theta'|\mathbf{z}^N)}f_{\theta'}(x) + (\mathbb{E}_{q(\theta'|\mathbf{z}^N)}f_{\theta'}(x))^2 + 2\mathbb{E}_{q(\theta|\mathbf{z}^N)}(f_\theta(x))^2 - \mathbb{E}_{q(\theta|\mathbf{z}^N)}(f_\theta(x))^2 \\
&= \mathbb{E}_{q(\theta|\mathbf{z}^N)}(f_\theta(x))^2 - (\mathbb{E}_{q(\theta'|\mathbf{z}^N)}f_{\theta'}(x))^2 \\
&= \mathbb{E}_{q(\theta|\mathbf{z}^N)}\|f_\theta(x) - \mathbb{E}_{q(\theta'|\mathbf{z}^N)}f_{\theta'}(x)\|^2 \\
&= \text{Var}_{\theta|\mathbf{z}^N}f_\theta(x).
\end{aligned} \tag{48}$$

This concludes the proof.

D.5 PROOF OF THEOREM 2

First, we prove Eq. (12). This is directly obtained by taking the expectation of $\nu(Y|x)$ for Lemma 1. We can prove Eq. (12) by the direct calculation. By definition

$$\begin{aligned}
\text{ER}^{(2)}(Y|x, \mathbf{z}^N, \theta^*) &:= R^{(2)}(Y|x, \mathbf{z}^N) - R^{(2)}(Y|x, \theta^*) \\
&= \mathbb{E}_{\nu(Y|x)q(\theta|\mathbf{z}^N)}\|Y - f_\theta(x)\|^2 - \mathbb{E}_{\nu(Y|x)}\|Y - \mathbb{E}_{\nu(Y|x)}[Y|x]\|^2 \\
&= -2f_{\theta^*}(x)\mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x) + \mathbb{E}_{q(\theta|\mathbf{z}^N)}[f_\theta(x)^2] + f_{\theta^*}(x)^2 \\
&= \|f_{\theta^*}(x) - \mathbb{E}_{q(\theta|\mathbf{z}^N)}f_\theta(x)\|^2 + \text{Var}_{\theta|\mathbf{z}^N}f_\theta(x) \\
&= \text{PER}^{(2)}(Y|x, \mathbf{z}^N) + \text{BER}^{(2)}(Y|x, \mathbf{z}^N),
\end{aligned} \tag{49}$$

where we used the relation $\mathbb{E}_{\nu(Y|x)}[Y|x] = f_{\theta^*}(x)$. By definition,

$$\mathbb{E}R^{(2)}(Y|x, \mathbf{z}^N, \theta^*) = R^{(2)}(Y|x, \mathbf{z}^N) - R^{(2)}(Y|x, \theta^*) \geq 0, \quad (50)$$

and for the squared loss for any action a , we have $l(y, a) \geq 0$. Combined these, we have

$$\mathbb{E}R^{(2)}(Y|x, \mathbf{z}^N, \theta^*) \leq R^{(2)}(Y|x, \mathbf{z}^N). \quad (51)$$

This concludes the proof of Eq. (12).

The unconditional relation is derived by using relations in Eqs. (34) and (35). Finally, we get Eq. (13) by the PAC-Bayesian bound Eq. (2).

D.6 DISCUSSION ABOUT $\mathcal{Y} = \mathbb{R}^d$

For $\mathcal{Y} = \mathbb{R}^d$, Lemma 1 and Theorem 1 hold since we can proceed the proof in the same way for $\mathcal{Y} = \mathbb{R}$. Thus, we can proceed the proof of Theorem 2 for $\mathcal{Y} = \mathbb{R}^d$ in the same way as $\mathcal{Y} = \mathbb{R}$. Thus Theorem 2 holds in $\mathcal{Y} = \mathbb{R}^d$.

As for Theorem 3, we consider $p(Y|x, \theta) = N(y|f_{\theta}(x), \text{diag}(v^2))$, where $v^2 \in \mathbb{R}^d$, where $\text{diag}(v^2)$ is the diagonal matrix with each entry is v_i^2 . Then Theorem 3 still holds.

D.7 PROOF OF THEOREM 3

First, we show Eq. (17). Recall the definition $L(y, x, \theta, \theta^*) = \ln p(y|x, \theta^*) - \ln p(y|x, \theta)$. We express this $L(Y, x, \theta) = L(y, x, \theta, \theta^*)$ for simplicity.

We use the following change-of-measure inequality, which is also known as the transportation lemma Boucheron et al. (2013); Xu & Raginsky (2020); Xu (2020).

Lemma 2. *Let W be a real-valued integrable random variable with probability distribution p . Let h be a convex and continuously differentiable function on a interval $(0, b]$ and assume $h(0) = h'(0) = 0$. Define for every $x \geq 0$, $h^*(x) = \sup_{0 \leq \rho < b} \{\rho x - h(\rho)\}$ and let for every $y \geq 0$, $h^{*-1}(y) := \sup\{x \in \mathbb{R} : h^*(x) \leq y\}$. Then if*

$$\ln \mathbb{E}_{p(W)} e^{\rho(W - \mathbb{E}_p(W)W)} \leq h(\rho), \quad (52)$$

is satisfied, for any probability distribution q , which is absolutely continuous with respect to p such that $\text{KL}(q|p) \leq \infty$, we have

$$\mathbb{E}_{q(W)} W - \mathbb{E}_{p(W)} W \leq h^{*-1}(\text{KL}(q|p)). \quad (53)$$

The proof of this lemma is shown in Xu & Raginsky (2020) as Theorem 4. Also, this lemma previously appeared in Boucheron et al. (2013) as Lemma 4.14.

When $h(\rho) = \rho^2 \sigma^2 / 2$ and $b = \infty$, this assumption is σ^2 sub-Gaussian property and $h^{*-1}(x) = \sqrt{2x}$.

Then under the assumption, from the Lemma 2, given θ and x and \mathbf{z}^N , we have

$$\mathbb{E}_{p(Y|x, \theta^*)} L(Y, x, \theta) - \mathbb{E}_{p(Y|x, \theta)} L(Y, x, \theta) \leq h^{*-1}(\text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))). \quad (54)$$

When we focus on a sub-Gaussian property, we have,

$$\mathbb{E}_{p(Y|x, \theta^*)} L(Y, x, \theta) - \mathbb{E}_{p(Y|x, \theta)} L(Y, x, \theta) \leq \sqrt{2\sigma^2 \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))}. \quad (55)$$

By taking the expectation $\mathbb{E}_{q(\theta|\mathbf{z}^N)}$, we have

$$\begin{aligned} & \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x, \theta^*)} L(Y, x, \theta) \\ & \leq \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x, \theta)} L(Y, x, \theta) + \mathbb{E}_{q(\theta|\mathbf{z}^N)} \sqrt{2\sigma^2 \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))} \\ & \leq \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x, \theta)} L(Y, x, \theta) + \\ & \quad \sqrt{\mathbb{E}_{q(\theta|\mathbf{z}^N)} 2\sigma^2} \sqrt{\mathbb{E}_{q(\theta|\mathbf{z}^N)} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))}, \end{aligned} \quad (56)$$

where we used the Holder inequality in the last line. From the definition of L ,

$$\begin{aligned} & \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x,\theta^*)} L(Y, x, \theta) \\ &= \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x,\theta^*)} [-\ln p(Y|x, \theta) + \ln p(Y|x, \theta^*)] \\ &\geq \mathbb{E}_{p(Y|x,\theta^*)} [-\ln \mathbb{E}_{q(\theta|\mathbf{z}^N)} p(Y|x, \theta) + \ln p(Y|x, \theta^*)] \\ &= \text{PER}^{\log}(Y|x, \mathbf{z}^N), \end{aligned} \quad (57)$$

where we used the assumption that $p(y|x, \theta^*) = \nu(y|x)$ and used the Jensen inequality for the logarithmic function.

By definition, we have

$$\mathbb{E}_{q(\theta|\mathbf{z}^N)} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta)) = \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*). \quad (58)$$

From the definition of L , and since we are considering log loss, we have

$$\begin{aligned} & -\mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x,\theta)} L(Y, x, \theta) \\ &= \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x,\theta)} [-\ln p(Y|x, \theta^*) + \ln p(Y|x, \theta)] \\ &\geq \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x,\theta)} [-\ln \mathbb{E}_{q(\theta|\mathbf{z}^N)} p(Y|x, \theta) + \ln p(Y|x, \theta)] \\ &= \text{BER}^{\log}(Y|x, \mathbf{z}^N). \end{aligned} \quad (59)$$

Combined these, we have

$$\text{PER}^{\log}(Y|x, \mathbf{z}^N) + \text{BER}^{\log}(Y|x, \mathbf{z}^N) \leq \sqrt{\mathbb{E}_{q(\theta|\mathbf{z}^N)} 2\sigma^2} \sqrt{\text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*)}. \quad (60)$$

Then by applying the assumption of the PAC-Bayesian bound, we can obtain Eq. (17)

As for Eq. (18), we can proceed the calculation by taking the expectation $\mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X=x)}$ instead of $\mathbb{E}_{q(\theta|\mathbf{z}^N)}$ after Eq. (55). Then we get the result of Eq. (18).

D.8 THE CASE OF GAUSSIAN LIKELIHOOD

Here we show that we can apply the theorem to the Gaussian likelihood. We define $p(y|x, \theta) = N(y|f_\theta(x), v^2)$. From the definition, we have

$$\mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X=x)} \mathbb{E}_{p(Y|x,\theta^*)} L(Y, x, \theta) = \text{PER}^{\log}(Y|X, \mathbf{Z}^N). \quad (61)$$

We also have

$$\mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X=x)} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta)) = \frac{1}{2v^2} \mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X=x)} \|f_\theta(x) - f_{\theta^*}(x)\|^2. \quad (62)$$

Moreover, by directly calculating the definition of the exponential moment, we have

$$\ln \mathbb{E}_{p(Y|x,\theta)} e^{\rho(L(Y,x,\theta) - \mathbb{E}_{p(Y|x,\theta)} L(Y,x,\theta))} = \frac{\rho^2}{2v^2} \|f_\theta(x) - f_{\theta^*}(x)\|^2. \quad (63)$$

Then by applying the Lemma 2, we have

$$\mathbb{E}_{p(Y|x,\theta^*)} L(Y, x, \theta) - \mathbb{E}_{p(Y|x,\theta)} L(Y, x, \theta) \leq \frac{1}{v^2} \|f_\theta(x) - f_{\theta^*}(x)\|^2. \quad (64)$$

This implies that

$$\mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X)} \sigma^2(X, \theta) \leq \mathbb{E}_{\nu(\mathbf{z}^N)q(\theta|\mathbf{z}^N)\nu(X)} \frac{1}{v^2} |f_\theta(X) - f_{\theta^*}(X)|^2 = 2\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*) \leq \frac{2C_1 \ln N}{N^\alpha} := \sigma_q^2 \quad (65)$$

Summarizing above, for the Gaussian likelihood, we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq 2\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*). \quad (66)$$

D.9 RELAXATION OF ASSUMPTION IN THEOREM 3

We assumed that $\nu(y|x) = p(y|x, \theta^*)$ holds for Theorem 3. We can relax this assumption for specific models. In the proof of Theorem 3, we used assumption $\nu(y|x) = p(y|x, \theta^*)$ for connecting the KL divergence with ER (excess risk) as

$$\mathbb{E}_{q(\theta|\mathbf{z}^N)} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta)) = \mathbb{E}_{q(\theta|\mathbf{z}^N)} \text{KL}(\nu(Y|x)|p(Y|x, \theta)) = \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*), \quad (67)$$

where the first equality comes from the assumption and the second equality comes from the definition of the excess risk of the log-loss.

It has been proved in Proposition 2 of Heide et al. (2020) that if the model $p(y|x, \theta)$ is the generalized linear model the assumption can be relaxed. To state that condition, we introduce the definitions of a GLM:

$$p(y|x, \theta) := \exp(x^\top \theta y - F(\theta) + r(y)). \quad (68)$$

Here, given $x \in \mathcal{X} \subset \mathbb{R}^d$ and the mean value parameter is given by $g^{-1}(x^\top \theta)$ where g is the link function. F is the normalizing constant and r is the reference measure. With this setting, if the GLM model satisfies

$$\mathbb{E}_{\nu(Y|x)}[Y|x] = g^{-1}(x^\top \theta). \quad (69)$$

Then we have

$$\text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta)) = \text{KL}(\nu(Y|x)|p(Y|x, \theta)). \quad (70)$$

This can be proved by the direct calculation. This relation implies that even if $\nu(y|x) \neq p(y|x, \theta^*)$ and Eq. (69) is satisfied, then we have

$$\mathbb{E}_{q(\theta|\mathbf{z}^N)} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta)) = \mathbb{E}_{q(\theta|\mathbf{z}^N)} \text{KL}(\nu(Y|x)|p(Y|x, \theta)) = \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*). \quad (71)$$

Then following the proof of Theorem 3 in Appendix D.7, Theorem 3 holds even for $\nu(y|x) \neq p(y|x, \theta^*)$.

The condition Eq. (69) implies that the mean function is well specified. This

D.10 ENTROPY CONVERGENCE RATE

Corollary 1. *Under the same assumption as Theorem 3, assume that $\ln p(y|x, \theta)$ satisfies the σ^2 sub-Gaussian property similarly to Assumption 1. Conditioned on (x, \mathbf{z}^N) , we have*

$$H[p^q(Y|x, \mathbf{z}^N)] \leq R^{\log}(Y|x, \mathbf{z}^N) + 2\sqrt{2\sigma_p^2 \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*)}, \quad (72)$$

and if the excess risk bound of the PAC-Bayesian theory Eq. (2) holds, we have

$$H[p^q(Y|X, \mathbf{Z}^N)] = H[p(Y|X, \theta^*)] + \mathcal{O}(\sqrt{\ln N/N^\alpha}). \quad (73)$$

Proof. If the log loss $-\ln p(Y|x, \theta)$ satisfies the σ^2 sub-Gaussian property similarly to Assumption 1, conditioned on $(x, \theta, \mathbf{z}^N)$, we have

$$\mathbb{E}_{p(Y|x, \theta^*)} \ln p(Y|x, \theta) - \mathbb{E}_{p(Y|x, \theta)} \ln p(Y|x, \theta) \leq \sqrt{2\sigma^2 \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))}. \quad (74)$$

Thus, by taking the expectation about $q(\theta|\mathbf{z}^N)$, we have

$$\mathbb{E}_{q(\theta|\mathbf{z}^N)} H[p(Y|x, \theta)] \leq R^{\log}(Y|x, \mathbf{z}^N) + \sqrt{\mathbb{E}_{q(\theta|\mathbf{z}^N)} 2\sigma^2 \sqrt{\text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*)}}. \quad (75)$$

From Eq. (17)

$$\text{PER}^{\log}(Y|x, \mathbf{z}^N) + \text{BER}^{\log}(Y|x, \mathbf{z}^N) \leq \sqrt{2\sigma_p^2 \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*)}, \quad (76)$$

and note that $\text{PER}^{\log}(Y|x, \mathbf{z}^N) \geq 0$ and $\text{BER}^{\log}(Y|x, \mathbf{z}^N) = H[p^q(Y|x, \mathbf{z}^N)] - \mathbb{E}_{q(\theta|\mathbf{z}^N)} H[p(Y|x, \theta)]$. Combined these inequalities, we have

$$H[p^q(Y|x, \mathbf{z}^N)] \leq R^{\log}(Y|x, \mathbf{z}^N) + 2\sqrt{2\mathbb{E}_{q(\theta|\mathbf{z}^N)}\sigma^2}\sqrt{\text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*)}. \quad (77)$$

Next, we take the expectation over $\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X)$ in Eq. (75) instead of $q(\theta|\mathbf{z}^N)$, we have

$$\begin{aligned} & H[p^q(Y|X, \mathbf{Z}^N)] - H[p(Y|X, \theta^*)] \\ & \leq R^{\log}(Y|X, \mathbf{Z}^N) - H[p(Y|X, \theta^*)] + 2\sqrt{2\mathbb{E}_{\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X)}\sigma^2}\sqrt{\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*)} \\ & = \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*) + 2\sqrt{2\mathbb{E}_{\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X)}\sigma^2}\sqrt{\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*)}. \end{aligned} \quad (78)$$

Then from the excess risk bound of the PAC-Bayesian theory Eq. (2), we have $\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*) = \mathcal{O}(\ln N/N^\alpha)$, we get the bound. \square

D.11 DISCUSSION ABOUT THE LOGISTIC REGRESSION

Here we discuss the relation among Bayesian excess risk, frequentist excess risk, and the generalization ability for the logistic regression. For logistic regression, we define the model as $p(Y = 1|x, \theta) = \text{sig}(\theta^\top \phi(x))$ where $\text{sig} = 1/(1 + e^{-x})$ is the sigmoid function and $\phi(x)$ is the feature vector.

We consider applying the following change-of-measure inequality; let W be a real-valued integrable random variable. If

$$\ln \mathbb{E}_{p(W)} e^{(W - \mathbb{E}W)} < \infty, \quad (79)$$

is satisfied, then, for any probability distribution Q , which is absolutely continuous with respect to P such that $\text{KL}(Q|P) \leq \infty$, we have

$$\mathbb{E}_{q(W)} W - \mathbb{E}_{p(W)} W \leq \ln \frac{\mathbb{E}_{p(W)} e^{\rho(W - \mathbb{E}_{p(W)} W)} + \text{KL}(q|p)}{\rho}. \quad (80)$$

Here we assume that $q := p(Y|x, \theta^*)dy$, $p := p(Y|x, \theta)dy$, and $W = L(Y, x, \theta, \theta^*)$ conditioned on $(x, \theta, \mathbf{z}^N)$. Here $L(y, x, \theta, \theta^*) = -\ln p(y|x, \theta) + \ln p(y|x, \theta^*)$.

Conditioned on $(x, \theta, \mathbf{z}^N)$, for $\rho \leq 1$, we have

$$\begin{aligned} & \ln \mathbb{E}_{p(Y|x, \theta)} e^{\rho(L(Y, x, \theta, \theta^*) - \mathbb{E}_{p(Y|x, \theta)} L(Y, x, \theta, \theta^*))} \\ & = \ln \mathbb{E}_{p(Y|x, \theta)} e^{\rho(-\ln p(Y|x, \theta) + \ln p(Y|x, \theta^*))} + \mathbb{E}_{p(Y|x, \theta)} (\ln p(Y|x, \theta) - \ln p(Y|x, \theta^*)) \\ & = \ln \int p(Y|x, \theta)^{1-\rho} p(Y|x, \theta^*)^\rho dy + \rho \text{KL}(p(Y|x, \theta)|p(Y|x, \theta^*)) \\ & = (\rho - 1) D_\rho(p(Y|x, \theta^*)|p(Y|x, \theta)) + \rho \text{KL}(p(Y|x, \theta)|p(Y|x, \theta^*)) \\ & \leq \rho \text{KL}(p(Y|x, \theta)|p(Y|x, \theta^*)), \end{aligned} \quad (81)$$

where we used the definition

$$D_\alpha(P|Q) := \frac{1}{\alpha - 1} \ln \int \left(\frac{dP}{dQ} \right)^{\alpha-1} dP \geq 0. \quad (82)$$

Under this definition, we have

$$\begin{aligned} & \text{KL}(p(Y|x, \theta)|p(Y|x, \theta^*)) \\ & \leq \sup_{x \in \mathcal{X}, \theta \in \Theta} (-\ln \min\{\text{sig}(\theta^\top \phi(x)), 1 - \text{sig}(\theta^\top \phi(x))\}) \text{TV}(p(Y|x, \theta)|p(Y|x, \theta^*)) \\ & \leq \sup_{x \in \mathcal{X}, \theta \in \Theta} (-\ln \min\{\text{sig}(\theta^\top \phi(x)), 1 - \text{sig}(\theta^\top \phi(x))\}) \sqrt{\frac{1}{2} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))} \\ & \leq \sup_{x \in \mathcal{X}, \theta \in \Theta} (-\ln \min\{\text{sig}(\theta^\top \phi(x)), 1 - \text{sig}(\theta^\top \phi(x))\}) \sqrt{\text{ER}^{\log}(Y|X, \mathbf{z}^N, \theta^*)}, \end{aligned} \quad (83)$$

where we used the Pinsker inequality and used the assumption that the model is well-specified. For simplicity, we express the coefficient as $\rho := \sup_{x \in \mathcal{X}, \theta \in \Theta} (-\ln \min\{\text{sig}(\theta^\top \phi(x)), 1 - \text{sig}(\theta^\top \phi(x))\})$. Then, by the transportation lemma using $\rho = 1$, we have

$$\text{PER}^{\log}(Y|x, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq \rho \sqrt{\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*)} + \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*). \quad (84)$$

Then by taking the expectation with respect to $\nu(\mathbf{Z}^N)\nu(X)$, we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq \rho \sqrt{\text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*)} + \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*), \quad (85)$$

where we used the Jensen inequality. For the logistic model, the PAC-Bayesian bound Eq. (2) holds (Alquier, 2021), we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq \rho \sqrt{C_1 \left(\frac{\ln N}{N^\alpha}\right)} + C_1 \left(\frac{\ln N}{N^\alpha}\right). \quad (86)$$

Since $\text{BER}^{\log}(Y|X, \mathbf{Z}^N) = I_\nu(\theta; Y|X, \mathbf{Z}^N)$, thus the mutual information converges $\mathcal{O}\left(\sqrt{\frac{\ln N}{N^\alpha}}\right)$.

Moreover using Collorary 1, we can derive the convergence of the entropy. Considering the same calculation in the proof of Collorary 1, we have

$$\mathbb{E}_{q(\theta|\mathbf{z}^N)} H[p(Y|x, \theta)] \leq R^{\log}(Y|x, \mathbf{z}^N) + \rho \sqrt{\text{ER}^{\log}(Y|X, \mathbf{z}^N, \theta^*)} + \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*). \quad (87)$$

Note that $\text{PER}^{\log}(Y|x, \mathbf{z}^N) \geq 0$ and $\text{BER}^{\log}(Y|x, \mathbf{z}^N) = H[p^q(Y|x, \mathbf{z}^N)] - \mathbb{E}_{q(\theta|\mathbf{z}^N)} H[p(Y|x, \theta)]$, we have

$$H[p^q(Y|x, \mathbf{z}^N)] \leq R^{\log}(Y|x, \mathbf{z}^N) + 2\rho \sqrt{\text{ER}^{\log}(Y|X, \mathbf{z}^N, \theta^*)} + 2\text{ER}^{\log}(Y|X, \mathbf{z}^N, \theta^*). \quad (88)$$

Thus, the entropy is bounded by the test loss. Then by taking the expectation with respect to $\nu(\mathbf{Z}^N)\nu(X)$, and using the PAC-Bayesian bound Eq. (2). Then we have

$$H[p(Y|X, \mathbf{Z}^N)] \leq H[p(Y|X, \theta^*)] + 2\rho \sqrt{C_1 \left(\frac{\ln N}{N^\alpha}\right)} + 3C_1 \left(\frac{\ln N}{N^\alpha}\right). \quad (89)$$

Thus, we have

$$H[p(Y|X, \mathbf{Z}^N)] \leq H[p(Y|X, \theta^*)] + \mathcal{O}\left(\sqrt{\frac{\ln N}{N^\alpha}}\right). \quad (90)$$

D.12 DISCUSSION OF ENTROPIC RISK

Before showing the relation between posterior variance and mutual information, we introduce an important lemma used in the analysis.

Lemma 3 (Lemma 1 in (Haussler & Opper, 1997)). *Let $P(w)$ be a measure on a set W and $Q(v)$ be a measure on a set V . For any real-valued function $u(w, v)$, we have*

$$-\int_V dQ(v) \ln \int_W dP(w) e^{u(w, v)} \leq -\ln \int_W dP(w) e^{\int_V dQ(v) u(w, v)}. \quad (91)$$

For completeness, we show the proof.

Proof. For any real valued functions u_1 and u_2 and $0 \leq \alpha \leq 1$, we have

$$\begin{aligned} \int_W dP(w) e^{\alpha u_1(w) + (1-\alpha)u_2(w)} &= \int_W dP(w) \left(e^{u_1(w)}\right)^\alpha \left(e^{u_2(w)}\right)^{1-\alpha} \\ &\leq \left(\int_W dP(w) e^{u_1(w)}\right)^\alpha \left(\int_W dP(w) e^{u_2(w)}\right)^{1-\alpha}, \end{aligned} \quad (92)$$

where we used Hölder's inequality. Taking logarithmic function, this shows that $\ln \int_W dP(w) e^{u(w, v)}$ is convex in u . Thus, the result follows by using the Jensen inequality. \square

Thus, this theorem is not restricted to probability distributions.

We first show the relation between posterior variance and the mutual information ($\text{BER}^{\log}(Y|x, \mathbf{z}^N)$).

Lemma 4. *For the Gaussian likelihood $p(y|x, \theta) = N(y|f_\theta(x), v^2)$, we have*

$$\text{BER}^{\log}(Y|x, \mathbf{z}^N) \leq \frac{\text{Var} f_\theta(x)}{v^2}. \quad (93)$$

Proof.

$$\begin{aligned} \text{BER}^{\log}(Y|x, \mathbf{z}^N) &= I_\nu(\theta; Y|x, \mathbf{z}^N) \\ &= \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x, \theta)} [-\ln \mathbb{E}_{q(\theta'|\mathbf{z}^N)} p(Y|x, \theta') + \ln p(Y|x, \theta)] \\ &= -\mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x, \theta)} \ln \mathbb{E}_{q(\theta'|\mathbf{z}^N)} e^{\ln p(Y|x, \theta') - \ln p(Y|x, \theta)}. \end{aligned} \quad (94)$$

Then, applying Lemma 3, we have

$$\begin{aligned} \text{BER}^{\log}(Y|x, \mathbf{z}^N) &\leq -\ln \mathbb{E}_{q(\theta'|\mathbf{z}^N)} e^{\mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{p(Y|x, \theta)} \ln p(Y|x, \theta') - \ln p(Y|x, \theta)} \\ &\leq -\ln \mathbb{E}_{q(\theta'|\mathbf{z}^N)} e^{-\frac{1}{2v^2} \mathbb{E}_{q(\theta|\mathbf{z}^N)} p(Y|x, \theta) (Y - f_{\theta'}(x))^2 - (Y - f_\theta(x))^2} \\ &\leq -\ln \mathbb{E}_{q(\theta'|\mathbf{z}^N)} e^{-\frac{1}{2v^2} \mathbb{E}_{q(\theta|\mathbf{z}^N)} (v^2 + f_\theta^2(x) - 2f_{\theta'}(x)f_\theta(x) + f_{\theta'}^2(x) - (v^2 + f_\theta^2(x) + f_\theta^2(x) - 2f_\theta^2(x)))} \\ &= -\ln \mathbb{E}_{q(\theta'|\mathbf{z}^N)} e^{-\frac{1}{2v^2} (\mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta^2(x) - 2\mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(x)f_{\theta'}(x) + f_{\theta'}^2(x))} \\ &\leq \frac{\text{Var} f_\theta(x)}{v^2}. \end{aligned} \quad (95)$$

□

Next we show Eq. (20). We use the transportation inequality Eq. (80). Similarly to the derivation of Eq. (66), we have

$$\begin{aligned} &-\mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{\nu(Y|x)} \ln \frac{p(Y|x, \theta)}{\nu(Y|x)} + \text{BER}^{\log}(Y|x, \mathbf{z}^N) \\ &\leq \mathbb{E}_{q(\theta|\mathbf{z}^N)} \text{KL}(\nu(Y|x)|p(Y|x, \theta)) + \frac{1}{2v^2} \mathbb{E}_{q(\theta|\mathbf{z}^N)} \|f_\theta(x) - f_{\theta^*}(x)\|^2. \end{aligned} \quad (96)$$

Thus, we have

$$\begin{aligned} &-\mathbb{E}_{\nu(Y|x)} \ln \mathbb{E}_{q(\theta|\mathbf{z}^N)} N(Y|f_\theta(x), v^2) + \mathbb{E}_{\nu(Y|x)} \ln \nu(Y|x) \\ &\leq \mathbb{E}_{q(\theta|\mathbf{z}^N)} \frac{1}{2v^2} \|f_\theta(x) - f_{\theta^*}(x)\|^2 - \mathbb{E}_{\nu(Y|x)} \mathbb{E}_{q(\theta|\mathbf{z}^N)} \ln N(Y|f_\theta(x), v^2) \\ &\quad - \text{BER}^{\log}(Y|x, \mathbf{z}^N) + \mathbb{E}_{\nu(Y|x)} \ln \nu(Y|x) \\ &\leq \mathbb{E}_{\nu(Y|x)q(\theta|\mathbf{z}^N)} \frac{1}{2v^2} \|Y - f_\theta(x)\|^2 - \mathbb{E}_{\nu(Y|x)} \frac{1}{2v^2} \|Y - f_{\theta^*}(x)\|^2 - \text{BER}^{\log}(Y|x, \mathbf{z}^N) \\ &\quad - \mathbb{E}_{\nu(Y|x)} \mathbb{E}_{q(\theta|\mathbf{z}^N)} \ln N(Y|f_\theta(x), v^2) + \mathbb{E}_{\nu(Y|x)} \ln \nu(Y|x). \end{aligned} \quad (97)$$

Then, we have

$$\begin{aligned} &-\mathbb{E}_{\nu(Y|x)} \ln \mathbb{E}_{q(\theta|\mathbf{z}^N)} N(Y|f_\theta(x), v^2) \\ &\leq \mathbb{E}_{\nu(Y|x)q(\theta|\mathbf{z}^N)} \frac{1}{2v^2} \|Y - f_\theta(x)\|^2 - \text{BER}^{\log}(Y|x, \mathbf{z}^N) - \mathbb{E}_{\nu(Y|x)} \mathbb{E}_{q(\theta|\mathbf{z}^N)} \ln N(Y|f_\theta(x), v^2) \\ &\leq \mathbb{E}_{\nu(Y|x)q(\theta|\mathbf{z}^N)} \frac{1}{v^2} \|Y - f_\theta(x)\|^2 + \frac{1}{2} \ln 2\pi v^2 - \text{BER}^{\log}(Y|x, \mathbf{z}^N) \\ &\leq \mathbb{E}_{\nu(Y|x)} \frac{1}{v^2} \|Y - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(x)\|^2 + \frac{1}{v^2} \text{Var} f_\theta(x) - \text{BER}^{\log}(Y|x, \mathbf{z}^N) + \frac{1}{2} \ln 2\pi v^2. \end{aligned} \quad (98)$$

This concludes the proof of Eq. (20).

Next, we discuss the entropic risk for the general log-likelihood other than the Gaussian distribution. Note that

$$\mathbb{E}_{\nu(Y|x)} \text{Ent}_{\alpha=1}^{\log}(Y, x) = -\mathbb{E}_{\nu(Y|x)} \ln \mathbb{E}_{q(\theta|\mathbf{z}^N)} p(Y|x, \theta). \quad (99)$$

Then from Eq. (18), we have

$$\text{PER}^{\log}(Y|x, \mathbf{z}^N) \leq \sqrt{2\sigma_p^2 \text{ER}^{\log}(Y|x, \mathbf{z}^N, \theta^*)} - \text{BER}^{\log}(Y|x, \mathbf{z}^N). \quad (100)$$

Then by using the Cauchy-Schwartz inequality and since we assumed that the model well-specified, we have

$$-\mathbb{E}_{\nu(Y|x)} \ln \mathbb{E}_{q(\theta|\mathbf{z}^N)} p(Y|x, \theta) \leq -\mathbb{E}_{\nu} \mathbb{E}_{q(\theta|\mathbf{z}^N)} \ln p(Y|x, \theta) - \text{BER}^{\log}(Y|x, \mathbf{z}^N) + \frac{\sigma_p^2}{2}. \quad (101)$$

Thus, the entropic risk has small regularization about BER.

E DISCUSSION OF SUB-EXPONENTIAL AND SUB-GAMMA ASSUMPTION

In the main paper, we have shown Theorem 3 and Collorary 1 when the sub-Gaussian property holds. Here we present the results for sub-exponential and sub-Gamma property. Conditioned on θ , x , and \mathbf{z}^N , by using Lemma 2 we have

$$\mathbb{E}_{p(Y|x, \theta^*)} L(Y, x, \theta) - \mathbb{E}_{p(Y|x, \theta)} L(Y, x, \theta) \leq h^{*-1}(\text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))). \quad (102)$$

Then by proceeding the calculation in the same way as the proof of Theorem 3, we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq \mathbb{E}_{\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X=x)} h^{*-1}(\text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))), \quad (103)$$

and since h^{*-1} is concave, we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq h^{*-1}(\mathbb{E}_{\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X=x)} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta))), \quad (104)$$

We need to derive h^{*-1} . For example in Lemma 2, if $h(\lambda) = \frac{\sigma^2 \lambda^2}{2}$ for $0 \leq \lambda \leq 1/b$, then we have

$$h^{*-1}(y) = \begin{cases} \sqrt{2\sigma^2 y} & \text{if } y \leq \frac{\sigma^2}{2b} \\ by + \frac{\sigma^2}{2b} & \text{otherwise.} \end{cases} \quad (105)$$

This is called as sub-exponential property. Thus, if $\mathbb{E}_{\nu(\mathbf{Z}^N)q(\theta|\mathbf{Z}^N)\nu(X=x)} \text{KL}(p(Y|x, \theta^*)|p(Y|x, \theta)) = \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*) \leq \frac{\sigma^2}{2b}$, for some constant b and σ^2 , we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq \sqrt{2\sigma^2 \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*)}, \quad (106)$$

and otherwise we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq b \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*) + \frac{\sigma^2}{2b}. \quad (107)$$

Next, in Lemma 2, if $h(\lambda) = \frac{\sigma^2 \lambda^2}{2(1-c|\lambda|)}$ for $0 \leq \lambda \leq 1/c$ and $c > 0$, then we have

$$h^{*-1}(y) = \sqrt{2\sigma^2 y} + cy. \quad (108)$$

This is called as sub-gamma property. If σ^2 and c are upper bounded by positive constants $\sigma_0^2 < \infty$ and $c < c_0 < \infty$, then we have

$$\text{PER}^{\log}(Y|X, \mathbf{Z}^N) + \text{BER}^{\log}(Y|X, \mathbf{Z}^N) \leq \sqrt{2\sigma_0^2 \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*)} + c_0 \text{ER}^{\log}(Y|X, \mathbf{Z}^N, \theta^*). \quad (109)$$

F DETAILED DESCRIPTION OF THE PROPOSED METHOD

First, we show the PAC-Bayesian bound for our proposed method. Following the high-probability bound of Theorem 4, given a distribution $\nu(Z)$, for any prior distribution $p(\theta)$ over Θ independent of

\mathbf{z}^N and for any $\xi \in (0, 1)$ and $c > 0$, with probability at least $1 - \xi$ over the choice of training data \mathbf{z}^N , for all probability distributions $q(\theta|\mathbf{z}^N)$ over Θ , we have

$$\begin{aligned} & \mathbb{E}_{\nu(Z)} \left[\frac{|Y - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(X)|^2}{2v^2} + \lambda \frac{\text{Var}_{\theta|\mathbf{z}^N} f_\theta(X)}{2v^2} \right] + \frac{\ln 2\pi v^2}{2} \\ & \leq \frac{1}{N} \sum_{i=1}^N \left[\frac{|\mathbb{E}_{q(\theta|\mathbf{z}^N)} y_i - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(x_i)|^2}{2v^2} + \lambda \frac{\text{Var}_{\theta|\mathbf{z}^N} f_\theta(x_i)}{2v^2} \right] + \frac{\ln 2\pi v^2}{2} + \frac{\text{KL}(q|p) + \frac{1}{2} \ln \frac{1}{\xi} + \frac{1}{2} \Omega_{p,\nu}(c, N)}{cN}, \end{aligned} \quad (110)$$

where

$$\Omega_{p,\nu}(c, N) := \ln \mathbb{E}_{p(\theta)p(\theta')} \mathbb{E}_{\nu(\mathbf{Z})^N} \exp[cN(\mathbb{E}_{\nu(Z)} L(Z, \theta, \theta') - \frac{1}{N} \sum_{n=1}^N L(Z_n, \theta, \theta'))], \quad (111)$$

$$L(z, \theta, \theta') := \mathbb{E}_{q(\theta|\mathbf{z}^N)} \frac{|y - f_\theta(x)|^2}{2v^2} + (\lambda - 1) \frac{\mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta^2(x) - \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{q(\theta'|\mathbf{z}^N)} f_\theta(x) f_{\theta'}(x)}{2v^2}. \quad (112)$$

Proof. The proof is similar to Masegosa et al. (2020). First note that

$$\begin{aligned} & \frac{|y - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(x)|^2}{2v^2} + \lambda \frac{\text{Var}_{\theta|\mathbf{z}^N} f_\theta(x)}{2v^2} \\ & = \mathbb{E}_{q(\theta|\mathbf{z}^N)} \frac{|y - f_\theta(x)|^2}{2v^2} + (\lambda - 1) \frac{\text{Var}_{\theta|\mathbf{z}^N} f_\theta(x)}{2v^2} \\ & = \mathbb{E}_{q(\theta|\mathbf{z}^N)} \frac{|y - f_\theta(x)|^2}{2v^2} + (\lambda - 1) \frac{\mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta^2(x) - \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{q(\theta'|\mathbf{z}^N)} f_\theta(x) f_{\theta'}(x)}{2v^2}. \end{aligned} \quad (113)$$

Based on this, we define the tandem loss as

$$L(z, \theta, \theta') := \mathbb{E}_{q(\theta|\mathbf{z}^N)} \frac{|y - f_\theta(x)|^2}{2v^2} + (\lambda - 1) \frac{\mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta^2(x) - \mathbb{E}_{q(\theta|\mathbf{z}^N)} \mathbb{E}_{q(\theta'|\mathbf{z}^N)} f_\theta(x) f_{\theta'}(x)}{2v^2}. \quad (114)$$

Then by considering the prior $p(\theta, \theta') = p(\theta)p(\theta')$, using Theorem 4, we have

$$\begin{aligned} \mathbb{E}_{\nu(Z)q(\theta|\mathbf{z}^N)q(\theta'|\mathbf{z}^N)} L(Z, \theta, \theta') & \leq \frac{1}{N} \sum_{n=1}^N \mathbb{E}_{q(\theta|\mathbf{z}^N)q(\theta'|\mathbf{z}^N)} L(z_n, \theta, \theta') \\ & \quad + \frac{\text{KL}(q(\theta|\mathbf{z}^N)q(\theta'|\mathbf{z}^N)|p(\theta)p(\theta')) + \ln \xi^{-1} + \Omega_{p,\nu}(c, N)}{cN}, \end{aligned} \quad (115)$$

where

$$\Omega_{p,\nu}(c, N) := \ln \mathbb{E}_{p(\theta)p(\theta')} \mathbb{E}_{\nu(\mathbf{Z})^N} \exp[cN(\mathbb{E}_{\nu(Z)} L(Z, \theta, \theta') - \frac{1}{N} \sum_{n=1}^N L(Z_n, \theta, \theta'))]. \quad (116)$$

Since $\text{KL}(q(\theta|\mathbf{z}^N)q(\theta'|\mathbf{z}^N)|p(\theta)p(\theta')) = 2\text{KL}(q(\theta|\mathbf{z}^N)|p(\theta))$, by setting $c=2c'$, we get the result. \square

Thus, the constant $\Omega_{p,\nu}$ depends only on the setting of the problem. We optimize the right-hand side of Eq. (110) as the objective function.

Next, we discuss the relation between rBER and standard VI. The objective function of standard VI is

$$-\mathbb{E}_{q(\theta|\mathbf{z}^N)} \ln N(y|f_\theta(x), v^2) = \frac{|y - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(x)|^2 + \text{Var}_{\theta|\mathbf{z}^N} f_\theta(x)}{2v^2} + \frac{1}{2} \ln 2\pi v^2. \quad (117)$$

Thus, we can interpret that the log loss of the Gaussian likelihood corresponds to the prediction risk and Bayesian excess risk. Since the prediction risk corresponds to the prediction performance, the

standard VI implicitly controls the prediction performance and the Bayesian excess risk. Our rBER can be regarded as

$$\text{rBER}(\lambda) = \frac{1}{2v^2} \frac{1}{N} \sum_{n=1}^N \left(\text{PR}^{(2)}(y_n|x_n, \mathbf{z}^N) + \lambda \text{BER}^{(2)}(y_n|x_n, \mathbf{z}^N) \right) + \frac{1}{2} \ln 2\pi v^2 + \frac{1}{N} \text{KL}(q|p). \quad (118)$$

Thus, BER has a flexible weight for regularizing the uncertainty. Numerically, when $\lambda = 0$, this corresponds to the setting where we simply optimize $\text{PR}^{(2)}(y|x, \mathbf{z}^N)$. This means we only consider the fitting performance. We numerically found that $\lambda = 0$ results in large uncertainty due to the lack of regularization. When $\lambda = 1$, we found that the uncertainty is underestimated.

Finally, we remark the relation between Bayesian excess risk and $\text{Var} f_\theta(x)$. Since we focus on the log loss, thus we can consider the following type of objective function.

$$\frac{1}{N} \sum_{n=1}^N \left[\frac{\mathbb{E}_Q |y_n - \mathbb{E}_{q(\theta|\mathbf{z}^N)} f_\theta(x_n)|^2}{2v^2} + \lambda \text{BER}^{\log}(y_n|x_n, \mathbf{z}^N) \right] + \frac{\ln 2\pi v^2}{2} + \frac{1}{N} \text{KL}(q|p), \quad (119)$$

where we use the Bayesian excess risk directly, instead of $\text{Var} f_\theta(x)$. Note that from Appendix D.12, $\text{BER}^{\log}(Y|x, \mathbf{z}^N) \leq \text{Var} f_\theta(x)/v^2$ holds for the Gaussian likelihood. Thus, Eq. (119) and our BER behaves in a similar way. From the numerical point of view, implementing $\text{Var} f_\theta(x)$ is easier than Eq. (119) since we simply calculate the variance of the prediction.

G NUMERICAL EXPERIMENTS

In this section, we describe the detailed settings of the experiments. We also present the additional experimental results.

G.1 PARTICLE VARIATIONAL INFERENCE

Since we applied our BER to particle variational inference (PVI), we briefly introduce the PVI and existing methods. In PVI, we use the empirical distribution $\rho(\theta) = \frac{1}{M} \sum_{i=1}^M \delta_{\theta_i}(\theta)$ as the posterior distribution. Here $\delta_{\theta_i}(\theta)$ is the Dirac distribution that has a mass at θ_i . We refer to these M samples as particles. PVI Liu & Wang (2016); Wang et al. (2019) approximates the posterior through iteratively updating the empirical distribution by interacting them with each other:

$$\theta_i^{\text{new}} \leftarrow \theta_i^{\text{old}} + \eta v_i(\{\theta_{i'}^{\text{old}}\}_{i'=1}^M), \quad (120)$$

where $v(\{\theta\})$ is the update direction. Basically, v is composed of the gradient term and the repulsion term to enhance the diversity of the posterior distribution since we are often interested in the multi-modal information of the posterior distribution. For the update direction about v , see f-SVGD in Wang et al. (2019) and VAR in Futami et al. (2021) for details.

We follow the approach in VAR in Futami et al. (2021). They proposed using the gradient of the PAC-Bayesian bound for the update direction v in Eq. (120). For example, VAR uses

$$v_i = \partial_i \mathcal{F}(\{\theta_i\}_{i=1}^N), \quad (121)$$

$$\mathcal{F}(\{\theta_i\}_{i=1}^N) := -\frac{1}{NM} \sum_{i=1}^M \sum_{n=1}^N [\ln p(y_n|x_n, \theta_i) + R(y_n, x_n)] + \frac{1}{N} \text{KL}(\rho_E|\pi), \quad (122)$$

where R is the repulsion term to enhance the diversity. See Futami et al. (2021) for details. Following their setting, we consider using the following update direction

$$v_i = \partial_i \text{rBER}(\lambda) \quad (123)$$

$$\text{rBER}(\lambda) = \frac{1}{N} \sum_{n=1}^N \frac{|y_n - \mathbb{E}_{\rho(\theta)} f_\theta(x_n)|^2}{2v'^2} + \lambda \frac{\text{Var} f_\theta(x_n)}{2v'^2} + \frac{\ln 2\pi v'^2}{2} + \frac{1}{N} \text{KL}(\rho(\theta)|p(\theta)), \quad (124)$$

where $\rho(\theta) = \frac{1}{M} \sum_{i=1}^M \delta_{\theta_i}(\theta)$. We optimize v' by gradient descent.

G.2 TOY DATA EXPERIMENTS AND DEPTH ESTIMATION

For these experiments, we used the implementation in the previous work Amini et al. (2020). For the toy data experiments, we used the Adam optimizer with the stepsize 0.0001 in the implementation of Amini et al. (2020). The number of ensembles is 5. We set other hyperparameters as the same as in Amini et al. (2020).

As for the depth estimation experiments, we used the same hyperparameter setting in Amini et al. (2020). Here, the number of ensembles is 5.

G.3 BNN REGRESSION FOR UCI DATASET

We used the same setting as the previous work Wang et al. (2019); Futami et al. (2021). We used the Adam optimizer with a learning rate of 0.004. We used a batch size of 100 and ran 500 epochs for the dataset size to be smaller than 1000. For a larger dataset, we used a batch size of 1000 and ran 3000 epochs.

To calculate the PICP, we first calculate the 95% prediction interval. We then calculate the number of the test data points included inside the prediction interval.

To calculate the MPIW, we calculated the mean of the prediction interval and normalized it by the maximum length of the test data point; $\max y_{\text{test}} - \min y_{\text{test}}$.

We show the additional results here. We show the result of PAC_E^2 and the negative log-likelihood. We also show the result of the α -divergence minimization. Following the previous work Li & Gal (2017), we considered the entropic loss for α -divergence minimization. In the definition of Li & Gal (2017), $\alpha = 0$ corresponds to the standard (exclusive) KL divergence, $\alpha = 0.5$ corresponds to the Hellinger divergence, and $\alpha = 1.0$ corresponds to the (inclusive) KL divergence, which is used in expectation propagation algorithm. We test on $\alpha = 0.5, 1.0$ and 2.0 . The results are shown in Table 3 to 6.

First, we found that α -divergence minimization show similar behaviors to f-SVGD in RMSE and negative test log-likelihood. However, we found that α -divergence minimization shows very large uncertainties since their PICP are much larger than 0.95 and their MPIW are larger than f-SVGD.

We found that PAC_E^2 shows the similar to $\text{BER}(0)$ measured in the negative log-likelihood, MPIW, and PICP. However, the prediction performance of PAC_E^2 in RMSE is significantly worse than $\text{BER}(0)$. This is because the objective function of PAC_E^2 is the negative log-likelihood of the predictive distribution; thus, the performance in RMSE is not guaranteed. On the other hand, the objective function of $\text{BER}(0)$ is based on the squared loss. Thus, it can show performance in RMSE.

Next, we evaluated how the RMSE, PICP, MPIW, and negative log-likelihood behave by changing λ in BER . We show the results in Fig.3 and 4. We confirmed that the prediction performance measured in RMSE does not depend on the choice of λ . On the other hand, other measures depend on λ significantly. The ideal PICP is 0.95. Thus, we should choose λ by cross-validation. We also found that even $\lambda = 1$. correspond to the standard VI. It underestimates the PICP.

G.4 CONTEXTUAL BANDIT TASKS

Here we explain the setting of the task. Our experiments follow the setting in Wang et al. (2019); Futami et al. (2021). Denote the context set as \mathcal{S} . For each time step t , an agent receives context $s_t \in \mathcal{S}$ from the environment. The agent choose action $a_t \in \{1, \dots, A\}$ based on the context s_t and get a reward $r_{a_t, t}$. We would like to minimize the pseudo regret

$$R_T = \max_{\phi: \mathcal{S} \rightarrow \{1, \dots, A\}} \mathbb{E} \left[\sum_{t=1}^T r_{g(s_t), t} - \sum_{t=1}^T r_{a_t, t} \right], \quad (125)$$

where ϕ maps the context to the action. We consider a prior $\mu_{s, i, 0}$ over a reward of context s and action i . Then, we update the prior to a posterior distribution using the observed reward. Following the previous work, we use Thompson sampling to select the action as

$$r_t \in \arg \max_{i \in \{1, \dots, K\}} \hat{r}_{i, t}, \quad \hat{r}_{i, t} \sim \mu_{s, i, t}. \quad (126)$$

Table 3: Benchmark results on test RMSE

Dataset	Avg. Test RMSE							
	f-SVGD	$\alpha = 0.5$	$\alpha = 1.0$	$\alpha = 2.0$	VAR	PAC_E^2	rBER(0.05)	f-SVGD
Concrete	4.33±0.8	4.51±0.8	4.67±0.7	4.98±0.6	4.30±0.7	5.49±0.5	4.47±0.6	4.48±0.7
Boston	2.54±0.50	2.81±0.88	2.87±0.80	2.98±0.90	2.53±0.50	4.41±0.45	2.53±0.50	2.53±0.51
Wine	0.61±0.04	0.61±0.04	0.61±0.04	0.61±0.03	0.61±0.04	1.02±0.11	0.64±0.04	0.63±0.02
Power	3.78±0.14	3.78±0.11	3.78±0.13	3.80±0.11	3.75±0.13	5.24±0.45	3.66±0.15	3.69±0.12
Yacht	0.64±0.28	0.56±0.26	0.88±0.34	0.99±0.68	0.60±0.28	0.71±0.41	0.75±0.41	0.78±0.48
Protein	3.98±0.54	4.05±0.13	3.97±0.04	4.02±0.09	3.92±0.05	7.96±0.10	3.83±0.10	3.85±0.05

Table 4: Benchmark results on test PICP and MPIW.

Dataset	Avg. Test PICP and MPIW in parenthesis				
	f-SVGD	VAR	PAC_E^2	rBER(0)	rBER(0.05)
Concrete	0.82±0.03 (0.13±0.00)	0.87±0.04 (0.16±0.01)	0.97±0.02 (0.57±0.04)	0.99±0.02 (0.50±0.04)	0.95±0.02 (0.25±0.02)
Boston	0.63±0.07 (0.10±0.02)	0.76±0.05 (0.14±0.01)	0.94±0.04 (0.40±0.04)	0.97±0.01 (0.33±0.04)	0.92±0.04 (0.22±0.02)
Wine	0.79±0.03 (0.32±0.05)	0.85±0.02 (0.39±0.06)	0.98±0.01 (1.06±0.01)	0.99±0.00 (1.61±0.00)	0.95±0.03 (0.32±0.15)
Power	0.43±0.01 (0.07±0.00)	0.82±0.01 (0.15±0.00)	0.99±0.00 (0.57±0.02)	0.99±0.01 (0.81±0.01)	0.96±0.01 (0.37±0.01)
Yacht	0.92±0.04 (0.02±0.01)	0.93±0.04 (0.04±0.01)	0.97±0.03 (0.07±0.00)	0.96±0.03 (0.10±0.01)	0.94±0.04 (0.08±0.01)
Protein	0.53±0.01 (0.24±0.01)	0.83±0.00 (0.58±0.01)	0.98±0.00 (1.44±0.06)	1.0±0.00 (5.04±0.01)	0.96±0.01 (0.86±0.00)

Table 5: Benchmark results on test PICP and MPIW for α -divergence minimization.

Dataset	Avg. Test PICP and MPIW in parenthesis		
	$\alpha = 0.5$	$\alpha = 1.0$	$\alpha = 2.0$
Concrete	0.97±0.01 (0.41±0.03)	0.99±0.01 (0.46±0.03)	0.99±0.01 (0.52±0.04)
Boston	0.98±0.01 (0.51±0.08)	0.99±0.01 (0.54±0.08)	0.99±0.01 (0.57±0.09)
Wine	0.94±0.02 (0.54±0.08)	0.95±0.02 (0.58±0.09)	0.96±0.01 (0.63±0.09)
Power	0.98±0.00 (0.45±0.01)	0.99±0.00 (0.47±0.01)	1.00±0.00 (0.50±0.01)
Yacht	0.99±0.01 (0.16±0.01)	1.00±0.00 (0.30±0.04)	1.00±0.00 (0.60±0.06)
Protein	0.99±0.00 (1.13±0.02)	0.99±0.00 (1.12±0.03)	0.99±0.01 (1.15±0.04)

Table 6: Benchmark results on negative test log-likelihood

Dataset	Avg. negative test log likelihood							
	f-SVGD	$\alpha = 0.5$	$\alpha = 1.0$	$\alpha = 2.0$	VAR	PAC_E^2	rBER(0)	rBER(0.05)
Concrete	-2.85±0.15	-2.79±0.17	-2.82±0.08	-2.95±0.05	-2.81±0.06	-3.16±0.03	-3.50±0.03	-3.06±0.05
Boston	-2.34±0.31	-2.39±0.16	-2.43±0.12	-2.47±0.09	-2.34±0.24	-2.61±0.08	-2.55±0.05	-2.38±0.16
Wine	-0.89±0.08	-0.87±0.07	-0.88±0.05	-0.93±0.03	-0.90±0.06	-1.26±0.02	-1.84±0.03	-1.08±0.03
Power	-2.75±0.03	-2.73±0.02	-2.74±0.01	-2.80±0.03	-2.80±0.03	-3.17±0.03	-3.95±0.04	-2.86±0.01
Yacht	-0.81±0.67	-0.92±0.27	-0.77±0.17	-1.58±0.17	-0.87±0.38	-0.81±0.11	-1.62±0.17	-1.46±0.30
Protein	-2.70±0.00	-2.98±0.02	-2.90±0.01	-2.95±0.02	-2.84±0.00	-3.30±0.00	-4.45±0.02	-2.94±0.00

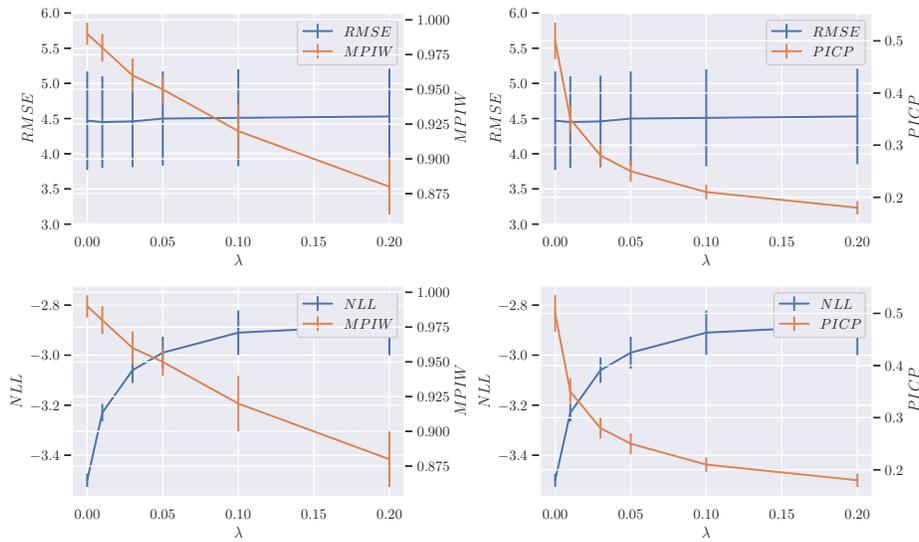


Figure 3: Concrete data in UCI dataset.

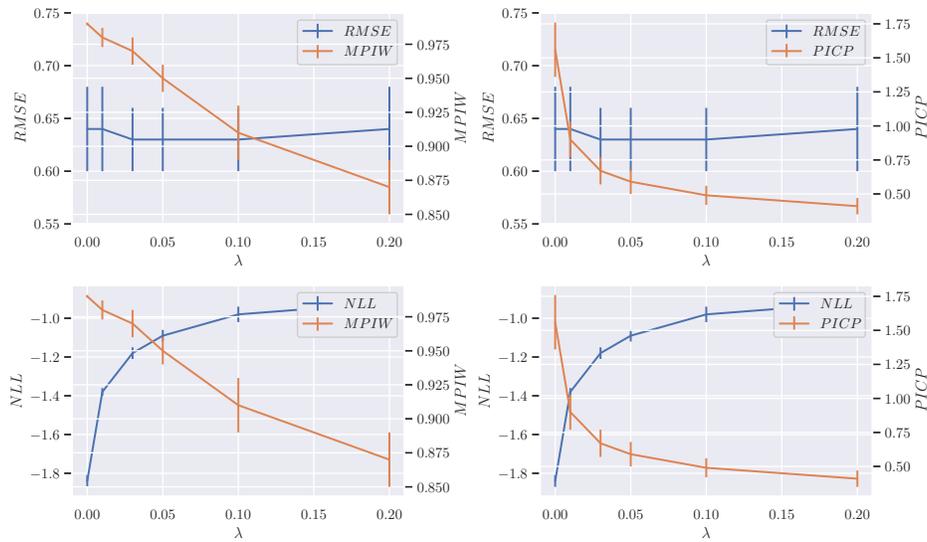


Figure 4: Wine data in UCI dataset.

We consider a neural network regression model following the previous work Wang et al. (2019); Futami et al. (2021), where the input is the context, and the output is the K -dimensional action. We place a prior distribution over the parameters of the network. We approximate the posterior distribution of the neural network parameters by PVI. All the hyperparameters are precisely the same as in the previous work Wang et al. (2019).

In Table 7, we additionally show performances of baseline methods, including Bayes by backpropagation (BBB), dropout, and black-box alpha divergence minimization ($BB\alpha$). These results are adapted from (Riquelme et al., 2018). As we can see, our proposed rBER(0.01) and rBER(0.05) outperformed these baseline methods.

Table 7: Cumulative regret relative to that of the uniform sampling with some existing methods.

Dataset	BBB	Dropout	BBα(= 1.0)	BBα(= 0.5)
Mushroom	0.036 \pm 0.002	0.056 \pm 0.012	0.543 \pm 0.000	0.539 \pm 0.000
Financial	0.23 \pm 0.013	0.17 \pm 0.007	0.40 \pm 0.007	0.70 \pm 0.035
Statlog	0.126 \pm 0.015	0.025 \pm 0.006	0.194 \pm 0.017	0.211 \pm 0.014
CoverType	0.582 \pm 0.022	0.306 \pm 0.002	0.394 \pm 0.005	0.605 \pm 0.000