

Calibration of classification models

Basics

- **Calibration**: how well the model's predictive probability aligns with the actual frequency of the true label.
- Binary classification $\mathcal{Y} = \{0, 1\}$
- ML models returns the predicted probability of $f_w : \mathcal{X} \rightarrow [0, 1]$
- Perfect calibration $\Pr(Y = 1 | f_w(X) = p) = p$ for all p a.s.

We evaluate this using **expected calibration error (ECE)**

$$ECE(f_w) := \mathbb{E}|\Pr[Y = 1 | f_w(X)] - f_w(X)|$$

After standard training, models are often not well calibrated.

- Q) How to estimate the ECE accurately using finite samples?
- Q) How to improve the calibration by post-processing the hypothesis?

Recalibration algorithm by PAC-Bayes

- A PAC-Bayesian generalization error analysis for the ECE.
- A recalibration algorithm based on Gaussian process derived by minimizing an upper bound on the PAC-Bayesian bound.

Generalization of ECE.

While the ECE is small on the training dataset, what can be said about

- (1) an unseen test dataset?
- (2) the population distribution?

PAC-Bayes bound (2)

$$\mathbb{E}_{V \sim \tilde{p}} |ECE(\eta_V \circ f_w) - \widehat{ECE}(\eta_V \circ f_w, S_{re})| \leq \frac{1+L}{B} + \frac{KL(\tilde{p} \| \tilde{\pi}) + B \log 2 + \log \frac{1}{\epsilon} + \frac{2\lambda^2}{n_{re}}}{\lambda}$$

Hypothesis after recalibration

Hypothesis before recalibration

Binning ECE

- We split the $[0,1]$ into B intervals $\mathcal{I} = \{I_1, \dots, I_B\}$
- Aggregate the values allocated into the same bins. (f and Y).

We use a **Gaussian process** as the recalibration model and optimize it in a manner analogous to **variational inference**.

M. Fujisawa* & F. Futami*. [PAC-Bayes Analysis for Recalibration in Classification](#). In International Conference on Machine Learning, 2025. (ICML 2025)

Estimating ECE for PU setting

- Constructing an estimator ECE using only **PU data**.
- Analyzing the bias and generalization error of the estimator.
- Clarifying the effect of class-prior estimation bias in ECE.

$$S_{PU} = \{X_m^{trU}\}_{m=1}^{n_U} \sim P(X)^{n_U} \cup \{X_m^{trP}\}_{m=1}^{n_P} \sim P(X|Y=1)^{n_P}$$

Proposed estimator: PU-ECE

$$ECE_{PU}(f, S_{PU}) := \sum_{b=1}^B \left| \frac{\pi_P}{n_P} \sum_{m=1}^{n_P} \mathbb{1}_{f(X_m^P) \in I_b} - \frac{1}{n_U} \sum_{m=1}^{n_U} \mathbb{1}_{f(X_m^U) \in I_b} \right|$$

$\pi_P = P(Y=1)$: class prior, $S_{PU} = (\{X_1^U, \dots, X_{n_U}^U\}, \{X_1^P, \dots, X_{n_P}^P\})$: PU evaluation data

Rewrite supervised ECE as

$$\sum_{b=1}^B \left| P(f(X) \in I_b) \mathbb{E}[f(X) | f(X) \in I_b] - \underbrace{P(Y=1 | f(X) \in I_b)}_{\parallel} \right|$$

- can be estimated from positive data $P(Y=1) P(f(X) \in I_b | Y=1)$
- can be estimated from unlabeled data $P(f(X) \in I_b)$
- is known or can be estimated

Under the same conditions as in PN learning

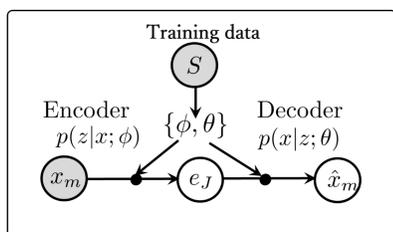
$$Bias_{tot}(f, S_{PU}) := |ECE_{PU}(f, S_{PU}) - TCE(f)| = o\left(\left(\frac{\pi_P^2}{n_P} + \frac{1}{n_U}\right)^{\frac{1}{2}}\right)$$

R. Kiryo, F. Futami, M. Sugiyama, [Estimating Expected Calibration Error for Positive-Unlabeled Learning](#). Transactions on Machine Learning Research, 2025.

Probabilistic models

Information theoretic analysis of VQ-VAE

- We propose a new **information-theoretic generalization analysis for VQ-VAEs** that elucidates the behavior of latent variables.
- The resulting **bound depends only on encoder overfitting and the complexity of the LVs, with no explicit dependence on the decoder**.
- We also analyze the generation performance.



- **Encoder-Decoder model**: Encoder network compresses the data into **latent variables (LVs), a.k.a representations**.
- Many learning algorithms use regularization over LVs.

min Training error + **Complexity of LVs**

Network parameters

Example

- Variational inference $KL(p(Z|X; \phi) \| p(Z))$
- Information bottleneck $I(Z; X)$

Q). Regularizing the LVs imply generalization?

Information-theoretic generalization analysis

- We introduced a new data-dependent prior distribution over LVs,

Generalization of the reconstruction loss: the squared-loss discrepancy between the reconstructed data and the original data.

$$gen(n, \mathcal{D}) \leq 3\Delta \sqrt{\frac{I(\tilde{J}; \mathbf{T} | \mathbf{e}, \phi, \tilde{X}) + I(\mathbf{e}, \phi; S)}{n}} + \frac{\Delta}{\sqrt{n}}$$

Complexity of **LVs**

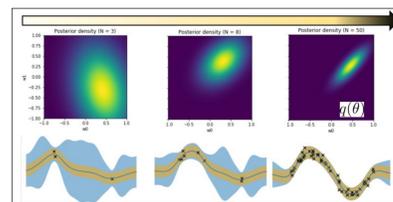
Complexity of **Encoder network**

- Controlling the complexity of **both LVs and the encoder network** is crucial, whereas the complexity of the decoder has minimal effect.

F. Futami* & M. Fujisawa*. [Information-theoretic Generalization Analysis for VQ-VAEs: A Role of Latent Variables](#). Neural Information Processing Systems, 2025. (NeurIPS 2025)

Evaluation of Epistemic uncertainty by PAC-Bayes

- We analyze common Bayesian uncertainty measures via PAC-Bayes.
- We show they lower-bound the excess risk and study their sample-size dependence.
- We propose a variational Bayesian method to control epistemic uncertainty.



Aleatoric uncertainty: Inherent noise in the data-generating process (fundamental difficulty of learning)

Epistemic uncertainty: Insufficient training data (when the model is well-specified)

$$Total\ risk = \underbrace{Bayes\ risk}_{Aleatoric} + \underbrace{Excess\ risk}_{Epistemic\ uncertainty}$$

Widely used measures of predictive uncertainty lower bounds the excess risk

- **the predictive mean** \Leftrightarrow **squared loss**
- **the MI between the prediction and model** \Leftrightarrow **Log loss**

$$BER^{\log}(Y|x, \mathbf{z}^N) \leq$$

$$2 \left(\underbrace{\sigma^2 KL(\nu(Y|x) \| p(Y|x, \theta^*))}_{Error\ due\ to\ model\ misspecification} + \underbrace{\sigma^2 ER^{\log}(Y|x, \mathbf{z}^N, \theta^*)}_{Error\ due\ to\ insufficient\ data} \right)^{\frac{1}{2}}$$

MI optimistically evaluates the excess risk, assuming that our predictive model is correct.

F. Futami. [Epistemic Uncertainty and Excess Risk in Variational Inference](#). In Artificial Intelligence and Statistics, 2025. (AISTATS 2025)

Convergence analysis of SVGD

- We present a method for analyzing convergence to the posterior in SVGD, an approximate Bayesian inference method.
- We propose a new approach for handling approximation error induced by projection.

M. Fujisawa* & F. Futami*. [On the Convergence of SVGD in KL divergence via Approximate gradient flow](#). Transactions on Machine Learning Research, 2025.