

A Note on the Chernoff Bound for Random Variables in the Unit Interval

Andrew Y. K. Foong¹, Wessel P. Bruinsma^{1, 2}, David R. Burt¹

¹University of Cambridge, ²Invenia Labs
 {ykf21, wpb23, drb62}@cam.ac.uk

May 2022

Abstract

The Chernoff bound is a well-known tool for obtaining a high probability bound on the expectation of a Bernoulli random variable in terms of its sample average. This bound is commonly used in statistical learning theory to upper bound the generalisation risk of a hypothesis in terms of its empirical risk on held-out data, for the case of a binary-valued loss function. However, the extension of this bound to the case of random variables taking values in the unit interval is less well known in the community. In this note we provide a proof of this extension for convenience and future reference.

1 Introduction

In statistical learning theory, one commonly considers a hypothesis space \mathcal{H} and a probability measure D over a space of datapoints \mathcal{Z} . Let $\ell: \mathcal{H} \times \mathcal{Z} \rightarrow \{0, 1\}$ be a loss function taking values in $\{0, 1\}$; *i.e.*, ℓ is a binary loss. It is often of interest to bound the generalisation risk $R(h) := \mathbb{E}_{Z \sim D}[\ell(h, Z)]$ of a hypothesis $h \in \mathcal{H}$. Such a bound can be established by computing the empirical mean of the loss on a dataset $S \sim D^N$ and using the following well-known theorem:

Theorem 1 (Chernoff bound for binary random variables; [Langford, 2005](#), Corollary 3.7). *Let X_1, \dots, X_n be i.i.d. random variables with $X_i \in \{0, 1\}$ and $\mathbb{E}[X_i] = p$. Let $\bar{X} := \frac{1}{n} \sum_{i=1}^n X_i$. Then, with probability at least $1 - \delta$,*

$$p \leq \text{kl}^{-1} \left(\bar{X} \left| \frac{1}{n} \log \frac{1}{\delta} \right. \right), \quad (1)$$

where $\text{kl}(q, p) := q \log \frac{q}{p} + (1 - q) \log \frac{1-q}{1-p}$ and $\text{kl}^{-1}(q | c) := \sup \{p \in [0, 1] : \text{kl}(q, p) \leq c\}$.

By choosing $X_i = \ell(h, Z_i)$ to be the loss of hypothesis h on the i^{th} datapoint Z_i in S , we immediately obtain a high-probability upper bound on the generalisation risk, assuming that h does not depend on S . (*I.e.*, if h is the result of training algorithm which uses training data, then S must be held-out data.)

Theorem 1 is well-known and, as stated, requires X_i to be Bernoulli random variables. However, if the loss ℓ takes values in the unit interval $[0, 1]$ rather than in $\{0, 1\}$, *i.e.* $\ell: \mathcal{H} \times \mathcal{Z} \rightarrow [0, 1]$, then we require an extension of theorem 1 to the case of X_i taking values in $[0, 1]$:

Theorem 2 (Chernoff bound for random variables in the unit interval). *Let X_1, \dots, X_n be i.i.d. random variables with $X_i \in [0, 1]$ and $\mathbb{E}[X_i] = p$. Then, using the same notation as in theorem 1, with probability at least $1 - \delta$,*

$$p \leq \text{kl}^{-1} \left(\bar{X} \left| \frac{1}{n} \log \frac{1}{\delta} \right. \right). \quad (2)$$

Although losses ℓ taking values in $[0, 1]$ are commonly encountered, theorem 2 is somewhat less well known than theorem 1. In this note, we provide a proof of theorem 2 for convenience and future reference.

2 Two-Sided Chernoff Bound

We first recapitulate some well-known bounds based on Hoeffding's extension of the Chernoff bound:

Lemma 2.1 (Hoeffding's extension). *Let X_1, \dots, X_n be i.i.d. random variables with $X_i \in [0, 1]$ and $\mathbb{E}[X_i] = p$. Then for any $t \in [0, p]$,*

$$\Pr(\bar{X} \leq p - t) \leq \exp(-n\text{kl}(p - t, p)), \quad (3)$$

and for any $t \in [0, 1 - p]$,

$$\Pr(\bar{X} \geq p + t) \leq \exp(-n\text{kl}(p + t, p)). \quad (4)$$

Proof. The proof is an application of the Chernoff method along with the observation that $z \mapsto e^{\lambda z}$ is convex, which allows us to control the moment-generating function in terms of the moment-generating function of a Bernoulli random variable. A detailed proof of the upper bound in eq. (3) is given in Theorem 5.1 of [Mulzer \(2018\)](#). Equation (4) follows by applying the change of variables $X_i \rightarrow 1 - X_i$ to eq. (3) (see Corollary 4.1 in [Mulzer \(2018\)](#) for an identical change of variables argument in the binary case). \square

We can use lemma 2.1 to obtain a *two-sided* bound on the mean p .

Theorem 3 (Two-sided Chernoff bound for random variables in the unit interval). *Let X_1, \dots, X_n be i.i.d. random variables with $X_i \in [0, 1]$ and $\mathbb{E}[X_i] = p$. Then, with probability at least $1 - \delta$,*

$$\text{kl}(\bar{X}, p) \leq \frac{1}{n} \log \frac{2}{\delta}. \quad (5)$$

Proof. For $c > 0$, let $\bar{\text{kl}}(p, c)$ be the unique real number in $(p, 1]$ such that $\text{kl}(\bar{\text{kl}}(p, c), p) = c$. Similarly, let $\underline{\text{kl}}(p, c)$ be the unique real number in $[0, p)$ such that $\text{kl}(\underline{\text{kl}}(p, c), p) = c$. Then, $\text{kl}(\bar{X}, p) \leq c$ if and only if $\underline{\text{kl}}(p, c) \leq \bar{X} \leq \bar{\text{kl}}(p, c)$. Hence,

$$\Pr\left(\text{kl}(\bar{X}, p) \leq \frac{1}{n} \log \frac{2}{\delta}\right) = \Pr\left(\underline{\text{kl}}\left(p, \frac{1}{n} \log \frac{2}{\delta}\right) \leq \bar{X} \leq \bar{\text{kl}}\left(p, \frac{1}{n} \log \frac{2}{\delta}\right)\right). \quad (6)$$

But, from eq. (3), we have

$$\Pr\left(\underline{\text{kl}}\left(p, \frac{1}{n} \log \frac{2}{\delta}\right) \leq \bar{X}\right) \geq 1 - \exp\left(-n\text{kl}\left(\underline{\text{kl}}\left(p, \frac{1}{n} \log \frac{2}{\delta}\right), p\right)\right) = 1 - \frac{\delta}{2}. \quad (7)$$

A symmetric argument shows that

$$\Pr\left(\bar{X} \leq \bar{\text{kl}}\left(p, \frac{1}{n} \log \frac{2}{\delta}\right)\right) \geq 1 - \frac{\delta}{2}. \quad (8)$$

Using a union bound therefore implies that

$$\Pr\left(\text{kl}(\bar{X}, p) \leq \frac{1}{n} \log \frac{2}{\delta}\right) \geq 1 - \delta, \quad (9)$$

which proves the theorem. \square

This *two-sided* bound on the mean p nearly implies the desired *one-sided* bound from theorem 2:

$$\text{kl}(\bar{X} | p) \leq \frac{1}{n} \log \frac{2}{\delta} \implies p \leq \text{kl}^{-1}\left(\bar{X} \mid \frac{1}{n} \log \frac{2}{\delta}\right). \quad (10)$$

Unfortunately, on the right-hand side, eq. (10) states $2/\delta$ instead of $1/\delta$.

3 One-Sided Chernoff Bound

We can tighten the $2/\delta$ in eq. (10) to obtain the $1/\delta$ from theorem 2 by directly proving a one-sided bound on the mean p .

Proof. The main ingredient of the proof of theorem 3 is the equivalence $\text{kl}(\bar{X}, p) \leq c$ if and only if $\underline{\text{kl}}(p, c) \leq \bar{X} \leq \overline{\text{kl}}(p, c)$. Since we now desire only an upper bound on p , $\text{kl}(\bar{X}, p) \leq c$ is stronger than we need. The key insight is to define a *one-sided* version of kl :

$$\text{kl}_m(q, p) = \begin{cases} \text{kl}(q, p), & q \leq p. \\ 0, & q > p. \end{cases} \quad (11)$$

An upper bound on p is then equivalent to an upper bound on $\text{kl}_m(\bar{X}, p)$:

$$p \leq \text{kl}^{-1} \left(\bar{X} \left| \frac{1}{n} \log \frac{1}{\delta} \right. \right) \iff \text{kl}_m(\bar{X}, p) \leq \frac{1}{n} \log \frac{1}{\delta}. \quad (12)$$

To see this, note that

$$\sup \left\{ p \in [0, 1] : \text{kl}(\bar{X}, p) \leq \frac{1}{n} \log \frac{1}{\delta} \right\} = \sup \left\{ p \in [0, 1] : \text{kl}_m(\bar{X}, p) \leq \frac{1}{n} \log \frac{1}{\delta} \right\} \quad (13)$$

because, in both suprema, $\bar{X} \leq p$ always. The same approach is taken in the proof of the one-sided bound for Bernoulli random variables in theorem 1, see the proof of Lemma 3.6 in Langford (2005).

Analogously to the proof of theorem 3, the main ingredient of this proof is the equivalence $\text{kl}_m(\bar{X}, p) \leq c$ if and only if $\underline{\text{kl}}(p, c) \leq \bar{X}$, so we conclude by the fact that the latter holds with probability at least $1 - \delta$:

$$\Pr \left(\text{kl}_m(\bar{X}, p) \leq \frac{1}{n} \log \frac{1}{\delta} \right) = \Pr \left(\underline{\text{kl}} \left(p, \frac{1}{n} \log \frac{1}{\delta} \right) \leq \bar{X} \right) \geq 1 - \delta. \quad (14)$$

□

References

- Langford, J. (2005). Tutorial on practical prediction theory for classification. *Journal of machine learning research*, 6(3).
- Mulzer, W. (2018). Five proofs of Chernoff’s bound with applications. *arXiv preprint arXiv:1801.03365*.