

---

# Adaptive Minimax Regret against Smooth Logarithmic Losses over High-Dimensional $\ell_1$ -Balls via Envelope Complexity

---

**Kohei Miyaguchi**  
IBM Research – Tokyo  
The University of Tokyo

**Kenji Yamanishi**  
The University of Tokyo

## Abstract

We develop a new theoretical framework, the *envelope complexity*, to analyze the minimax regret with logarithmic loss functions. Within the framework, we derive a Bayesian predictor that adaptively achieves the minimax regret over high-dimensional  $\ell_1$ -balls within a factor of two. The prior is newly derived for achieving the minimax regret and called the *spike-and-tails (ST) prior* as it looks like. The resulting regret bound is so simple that it is completely determined with the smoothness of the loss function and the radius of the balls except with logarithmic factors, and it has a generalized form of existing regret/risk bounds.

## 1 Introduction

As a notion of complexity of predictive models (sets of predictors), *minimax regret* has been considered in the literature of online learning (Cesa-Bianchi and Lugosi, 2006) and the minimum description length (MDL) principle (Rissanen, 1978; Grünwald, 2007). The minimax regret of a model  $\mathcal{H}$  is given by

$$\text{REG}^*(\mathcal{H}) = \inf_{\hat{h} \in \hat{\mathcal{H}}} \sup_{X \in \mathcal{X}} \left\{ f_X(\hat{h}) - \inf_{h \in \mathcal{H}} f_X(h) \right\}, \quad (1)$$

where  $f_X(h)$  denotes the loss of the prediction over data  $X$  made by  $h$ ,  $\hat{\mathcal{H}}$  denotes the feasible predictions and  $\mathcal{X}$  is the space of data. Here, the data may consist of a sequence of datum,  $X = X^n = (X_1, \dots, X_n)$ , and the loss maybe additive,  $f_X(h) = \sum_{i=1}^n f_{X_i}(h)$ , but we keep them implicit for generality. The minimax regret is a general complexity measure in the sense that it

is defined without any assumptions on the generation process of  $X$ . For instance, one can bound statistical risks with  $\text{REG}^*(\mathcal{H})$  regardless of the distribution of data (Littlestone, 1989; Cesa-Bianchi et al., 2004; Cesa-Bianchi and Gentile, 2008). Therefore, bounding the minimax regret and constructing the corresponding predictor  $\hat{h}$  is important to make a good and robust prediction.

We consider that  $\mathcal{H}$  is parametrized by a real-valued vector  $\theta \in \mathbb{R}^d$ ,  $\mathcal{H} = \{h_\theta \mid \gamma(\theta) \leq B, \theta \in \mathbb{R}^d\}$ , where  $\gamma(\theta)$  denotes a radius function such as norms of  $\theta$ . Thus, we may consider the luckiness minimax regret (Grünwald, 2007),

$$\text{LREG}^*(\gamma) = \inf_{\hat{h} \in \hat{\mathcal{H}}} \sup_{X \in \mathcal{X}} \left\{ f_X(\hat{h}) - \inf_{\theta \in \mathbb{R}^d} [f_X(\theta) + \gamma(\theta)] \right\}, \quad (2)$$

instead of the original minimax regret. Here, we abuse the notation  $f_X(\theta) = f_X(h_\theta)$ . There are at least three reasons for adopting this formulation. Firstly, as we do not assume the underlying distribution of  $X$ , it may be plausible to pose a soft restriction as in (2) rather than the hard restriction in (1). Secondly, it is straightforwardly shown that the luckiness minimax regret bounds above the minimax regret,  $\text{REG}^*(\mathcal{H}) \leq \text{LREG}^*(\lambda\gamma) + \lambda B$  for all  $\lambda \geq 0$ . Thus, it is often sufficient to bound  $\text{LREG}^*(\gamma)$  for bounding  $\text{REG}^*(\mathcal{H})$ . Finally, the luckiness minimax regret is including the original minimax regret as a special case such that  $\gamma(\theta) = 0$  if  $\theta \in \mathcal{H}$  and  $\gamma(\theta) = \infty$  otherwise. Therefore, we may avoid possible computational difficulties of the minimax regret by choosing the penalty  $\gamma$  carefully.

That being said, the closed-form expression of the exact (luckiness) minimax regret is even intractable except with few special cases (e.g., Shtar'kov (1987); Koolen et al. (2014)).

However, if we focus on information-theoretic settings, i.e., the model  $\mathcal{H}$  is a set of probabilistic distributions, everything becomes explicit. Now, let predictors be sub-probability distributions  $P(\cdot \mid \theta)$  and adopt the logarithmic loss function  $f_X(\theta) = -\ln \frac{dP}{d\nu}(X \mid \theta)$

with respect to an appropriate base measure  $\nu$  such as counting and Lebesgue measures. Note that a number of important practical problems such as logistic regression and data compression can be handled with this framework. With the logarithmic loss, the closed form of the luckiness minimax regret is given by Shtar'kov (1987); Grünwald (2007) as

$$\text{LREG}^*(\gamma) = \ln \int e^{-m(f_X + \gamma)} \nu(dX) \stackrel{\text{def}}{=} S(\gamma), \quad (3)$$

where  $m$  denotes the minimum operator given by  $m(f) = \inf_{\theta \in \mathbb{R}^d} f(\theta)$ . We refer to the right-hand side value as the *Shtarkov complexity*. Moreover, when all the distributions in  $\mathcal{H}$  are i.i.d. regular distributions of  $n$ -sequences  $X = (X_1, \dots, X_n)$ , under some regularity conditions, the celebrated asymptotic formula (Rissanen, 1996; Grünwald, 2007) is given as

$$S(\gamma) = \frac{d}{2} \ln \frac{n}{2\pi} + \int \sqrt{\det I(\theta)} e^{-\gamma(\theta)} d\theta + o(1), \quad (4)$$

where  $I(\theta)$  is the Fisher information matrix and  $o(1) \rightarrow 0$  as  $n \rightarrow \infty$ . More importantly, although the exact minimax-regret predictor achieving  $S(\gamma)$  is still intractable, the asymptotic formula implies that it is asymptotically achieved with the Bayesian predictor associated with the *tilted Jeffreys prior*  $\pi(d\theta) \propto \sqrt{\det I(\theta)} e^{-\gamma(\theta)} d\theta$ .

Here, our research questions are as follows: First, **(Q1)** *How can we evaluate  $S(\gamma)$  in modern high-dimensional contexts?* In particular, the asymptotic formula (4) does not withstand high-dimensional learning problems where  $d$  increases as  $n \rightarrow \infty$ . The exact evaluation of the Shtarkov complexity (3), on the other hand, is often intractable due to the minimum operator inside the integral. Second, **(Q2)** *How can we achieve the minimax regret with computationally feasible predictors?* It is important to provide the counterpart of the tilted Jeffreys prior in order to make actual predictions.

Regarding the above questions, our contribution is summarized as follows:

- We introduce the *envelope complexity*, a non-asymptotic approximation of the Shtarkov complexity  $S(\gamma)$  that allows us systematic computation of its upper bounds and predictors achieving these bounds. In particular, we show that the regret of the predictor is characterized with the smoothness.
- We demonstrate its usefulness by giving a Bayesian predictor that adaptively achieves the minimax regret within a factor of two over any high-dimensional smooth models under  $\ell_1$ -constraints  $\|\theta\|_1 \leq B$ .

The rest of the paper is organized as follows: In Section 2, we introduce the notion of Bayesian minimax regret as an approximation of the minimax regret within the ‘feasible’ set of predictors. We then develop a complexity measure called *envelope complexity* in Section 3 as a mathematical abstraction of the Bayesian minimax regret. We also present a collection of techniques for bounding the envelope complexity that is useful to bound the Shtarkov complexity and construct low-regret predictors. In Section 4, we utilize the envelope complexity to construct a near-minimax Bayesian predictor under  $\ell_1$ -penalization, namely the spike-and-tails (ST) prior. We also show that it achieves the minimax rate over  $\mathcal{H} = \{\theta \in \mathbb{R}^d \mid \|\theta\|_1 \leq B\}$  under high-dimensional asymptotics. In Section 5, we demonstrate numerical experiments to visualize our theoretical results. The discussion on these results in comparison to the existing studies is given in Section 6. Finally, we conclude the paper in Section 7.

## 2 Bayesian Minimax Regret

The minimax regret with logarithmic loss is given by the Shtarkov complexity  $S(\gamma)$ . The computation of the Shtarkov complexity is often intractable if we consider practical models such as deep neural networks. This is because the landscapes of loss functions  $f \in \mathcal{F}$  are complex as the models are, and hence their minimums  $m(f) = \inf_{\theta} f(\theta)$  and the complexity, which is an integral over the function of  $m(f)$ , are not tractable. Moreover, computations of the optimal predictor  $h^*$  are still often intractable even if  $S(\gamma)$  are given. For instance, the minimax-regret prediction for Bernoulli models over  $n$  outcomes takes  $O(n2^n)$  time. Of course there exist some special cases for which closed forms of  $\hat{h}$  are given. However, so far they are limited to exponential families.

One cause of this issue is that we seek for the best predictor  $\hat{h}$  among all the possible predictors  $\hat{\mathcal{H}}$ , i.e., all probability distributions. This is too general that it maybe not possible to compute  $\hat{h}$  nor  $\text{REG}^*(\gamma)$ . To avoid this difficulty, we narrow the set of feasible predictors  $\hat{\mathcal{H}}$  to the Bayesian predictors. Let  $w \in \mathcal{M}_+(\mathbb{R}^d)$  be a positive measure over  $\mathbb{R}^d$ , which we may refer to as *pre-prior*, and let  $\bar{h}_w$  be the Bayesian predictor associated with the prior  $\pi(d\theta) \propto e^{-\gamma(\theta)} w(d\theta)$ . Then we have

$$f_X(\bar{h}_w) = \ln \frac{w[e^{-\gamma}]}{w[e^{-f_X - \gamma}]} \stackrel{\text{def}}{=} f_X(w), \quad (5)$$

where  $w[\cdot]$  denotes the integral operation with respect to  $w(d\theta)$ . Now, we consider the Bayesian (luckiness)

minimax regret given by

$$\begin{aligned} \text{LREG}^{\text{Bayes}}(\gamma) &\stackrel{\text{def}}{=} \inf_{w \in \mathcal{M}_+(\mathbb{R}^d)} \text{LREG}(w|\gamma), \\ \text{LREG}(w|\gamma) &\stackrel{\text{def}}{=} \sup_{X \in \mathcal{X}} \{f_X(w) - m(f_X + \gamma)\}. \end{aligned}$$

One advantage of considering the Bayesian minimax regret is that, given a measure  $w$ , one can compute  $\bar{h}_w$  analytically or numerically utilizing techniques developed in the literature of Bayesian inference. In particular, a number of sophisticated variants of Monte Carlo Markov chain (MCMC) methods such as the stochastic gradient Langevin Dynamics (Welling and Teh, 2011) are developed for sampling  $\theta$  from complex posteriors.

Note that there does exist a case where the Bayesian minimax regret strictly differs from the minimax regret. See Barron et al. (2014) for example. It implies that narrowing the range of predictors to Bayesian may worsen the achievable worst-case regret. However, as we will show shortly, the gap between these minimax regrets can be controlled in some cases.

### 3 Envelope Complexity

We have introduced the Bayesian minimax regret  $\text{LREG}^{\text{Bayes}}(\gamma)$ . In this section, we present a set representation of Bayesian minimax regret, namely the *envelope complexity*  $C(\gamma, \mathcal{F})$ . Then, we show that the Shtarkov complexity is bounded by the envelope complexity and the envelope complexity can be easily bounded even if the models are complex.

#### 3.1 Set Representation of Bayesian Minimax Regret

The envelope complexity is a simple mathematical abstraction of Bayesian minimax regret and gives a fundamental basis for systematic computation of upper bounds on the (Bayesian) minimax regret. Let  $\mathcal{F}$  be a set of continuous functions  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  which is not necessarily logarithmic. Define the Bayesian envelope of  $\mathcal{F}$  as

$$\mathcal{E}(\mathcal{F}) \stackrel{\text{def}}{=} \left\{ w \in \mathcal{M}_+(\mathbb{R}^d) \mid \forall f \in \mathcal{F}, w[e^{-f+m(f)}] \geq 1 \right\},$$

and define the envelope complexity as

$$C(\gamma, \mathcal{F}) \stackrel{\text{def}}{=} \inf_{w \in \mathcal{E}(\mathcal{F})} \ln w[e^{-\gamma}].$$

Then, the envelope complexity characterizes Bayesian minimax regret.

**Theorem 1 (Set representation)** *Let  $\mathcal{F} = \{f_X + \gamma \mid X \in \mathcal{X}\}$ . Then, all measures in the*

*envelope  $w \in \mathcal{E}(\mathcal{F})$  satisfies that*

$$\text{LREG}(w|\gamma) \leq \ln w[e^{-\gamma}]. \quad (6)$$

*Moreover, we have*

$$\text{LREG}^{\text{Bayes}}(\gamma) = C(\gamma, \mathcal{F}).$$

**Proof** Let  $c(w) = \inf_{f \in \mathcal{F}} w[e^{-f+m(f)}]$ . Observe that

$$\begin{aligned} \ln \frac{w[e^{-\gamma}]}{c(w)} &= \sup_{f \in \mathcal{F}} \left\{ \ln \frac{w[e^{-\gamma}]}{w[e^{-f}]} - m(f) \right\} \\ &= \sup_{X \in \mathcal{X}} \left\{ \ln \frac{w[e^{-\gamma}]}{w[e^{-f_X - \gamma}]} - m(f_X + \gamma) \right\} \\ &\quad (f = f_X + \gamma) \\ &= \text{LREG}(w|\gamma). \\ &\quad (\because (5)) \end{aligned}$$

Then, since  $c(w) \geq 1$  for all  $w \in \mathcal{E}(\mathcal{F})$ , we have the first inequality.

Note that  $\bar{w} = w/c(w) \in \mathcal{E}(\mathcal{F})$  for any  $w \in \mathcal{M}_+(\mathbb{R}^d)$ , and  $\bar{w}[e^{-\gamma}] \leq w[e^{-\gamma}]$  whenever  $w \in \mathcal{E}(\mathcal{F})$ . Then we have

$$\begin{aligned} C(\gamma, \mathcal{F}) &= \inf_{w \in \mathcal{M}_+(\mathbb{R}^d)} \ln \frac{w[e^{-\gamma}]}{c(w)} \\ &= \inf_{w \in \mathcal{M}_+(\mathbb{R}^d)} \text{LREG}(w|\gamma) \\ &= \text{LREG}^{\text{Bayes}}(\gamma), \end{aligned}$$

yielding the second equality. This completes the proof.  $\blacksquare$

Whereas the first inequality (6) is helpful in construction of low-regret Bayesian predictors, the second equality shows that the envelope complexity is equivalent to the Bayesian minimax regret. Based on this result, we also present upper bounds of the Shtarkov complexity that are useful to estimate the complexity of hypothesis classes.

**Theorem 2 (Bounds on Shtarkov complexity)** *Let  $\mathcal{F} = \{f_X + \gamma \mid X \in \mathcal{X}\}$  where  $f_X$  is logarithmic. Then, for all  $w \in \mathcal{E}(\mathcal{F})$ , we have*

$$S(\gamma) \leq C(\gamma, \mathcal{F}) \leq \ln w[e^{-\gamma}].$$

**Proof** The first inequality follows from that the envelope minimax regret is no less than the minimax regret, as the range of infimum is shrunk from  $\hat{\mathcal{H}}$  to the Bayes class  $\{\bar{h}_w\}$ . The second inequality is seen by that the definition of the envelope complexity. This completes the proof.  $\blacksquare$

### 3.2 Useful Lemmas for Evaluating Envelope Complexity

Next, we show several lemmas that highlight the computational advantage of the Bayesian envelope and the envelope complexity. We start to show that the envelopes are easily evaluated with the surrogate relation. We say a function  $g$  is *surrogate* of another function  $f$  if and only if  $f - m(f) \leq g - m(g)$ , which is denoted by  $f \preceq g$ . Moreover, if there is one-to-one correspondence between  $g \in \mathcal{G}$  and  $f \in \mathcal{F}$  such that  $f \preceq g$ , then we may write  $\mathcal{F} \preceq \mathcal{G}$ .

**Lemma 3 (Monotonicity)** *Let  $\mathcal{F} \preceq \mathcal{G}' \subset \mathcal{G}$ . Then we have*

$$\mathcal{E}(\mathcal{F}) \supset \mathcal{E}(\mathcal{G}')$$

and therefore

$$C(\gamma, \mathcal{F}) \leq C(\gamma, \mathcal{G}').$$

**Proof** Note that  $e^{-f+m(f)} \geq e^{-g+m(g)}$  if  $f \preceq g$ , which means  $\mathcal{E}(\mathcal{F}) \supset \mathcal{E}(\mathcal{G}')$ . Also, as increasing the argument from  $\mathcal{G}'$  to  $\mathcal{G}$  just strengthen the predicate of the envelope, we have  $\mathcal{E}(\mathcal{G}') \supset \mathcal{E}(\mathcal{G})$ . Therefore, we have

$$\begin{aligned} C(\gamma, \mathcal{F}) &= \inf_{w \in \mathcal{E}(\mathcal{F})} \ln w [e^{-\gamma}] \\ &\leq \inf_{w \in \mathcal{E}(\mathcal{G}')} \ln w [e^{-\gamma}] \quad \mathcal{E}(\mathcal{F}) \supset \mathcal{E}(\mathcal{G}') \\ &\leq \inf_{w \in \mathcal{E}(\mathcal{G})} \ln w [e^{-\gamma}] \quad \mathcal{E}(\mathcal{G}') \supset \mathcal{E}(\mathcal{G}) \\ &= C(\gamma, \mathcal{G}). \end{aligned}$$

■

This is especially useful when the loss functions  $\mathcal{F}$  are complex but there exist simple surrogates  $\mathcal{G}$ . Consider any models such that the landscapes of the associated loss functions  $f \in \mathcal{F}$  are not fully understood and the evaluation of  $m(f)$  is expensive. It is impossible to check if  $w$  is in the envelope,  $w \in \mathcal{E}(\mathcal{F})$ , and therefore neither Theorem 1 nor 2 can be used directly. However, even in such cases, one can possibly find a surrogate class  $\mathcal{G}$  of  $\mathcal{F}$ . If the surrogate  $\mathcal{G}$  is simple enough for checking if  $w \in \mathcal{E}(\mathcal{G})$ , it is possible to bound the regret and the complexity utilizing Lemma 3.

In what follows, we consider the specific instance of the surrogate relation based on the smoothness. A function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  is *L-upper smooth* if and only if, for all  $\theta, \theta_0 \in \mathbb{R}^d$ , there exists  $g \in \mathbb{R}^d$  such that

$$f(\theta) \leq f(\theta_0) + g^\top (\theta - \theta_0) + \frac{L}{2} \|\theta - \theta_0\|_2^2. \quad (7)$$

Note that the upper smoothness is weaker than Lipschitz smoothness. Now, if  $f$  is *L-upper smooth* and has

at least one minima  $\theta_0 \in \arg m(f)$ , we can construct a simple quadratic surrogate of  $f$ ,  $\theta \mapsto \frac{L}{2} \|\theta - \theta_0\|_2^2$  ( $\succeq f$ ).

Motivated by the smoothness assumption, below we present more specific bounds for quadratic functions. Let  $\mathcal{Q}$  be the set of all quadratic functions with curvature one, defined as  $\mathcal{Q} = \{\theta \mapsto \frac{1}{2} \|\theta - u\|_2^2 \mid u \in \mathbb{R}^d\}$ . Moreover, for all sets of loss functions  $\mathcal{F}$  and penalty functions  $\gamma : \mathbb{R} \rightarrow \overline{\mathbb{R}}$ , we write  $\mathcal{F}_\gamma = \mathcal{F} + \gamma = \{f + \gamma \mid f \in \mathcal{F}\}$ . Then, the envelope of  $\mathcal{F}_\gamma$  is evaluated with that of  $\mathcal{Q}_\gamma$ .

**Lemma 4 (Smoothness-based bounds)** *Suppose that all  $f \in \mathcal{F}$  are L-upper smooth. Let  $\varphi(\theta) = \sqrt{L}^{-1} \theta$  be the scaling function. Then we have*

$$\mathcal{E}(\mathcal{Q}_{\gamma \circ \varphi}) \circ \varphi^{-1} \subset \mathcal{E}(\mathcal{F}_\gamma),$$

and moreover,

$$C(\gamma, \mathcal{F}_\gamma) \leq C(\gamma \circ \varphi, \mathcal{Q}_{\gamma \circ \varphi}).$$

**Proof** Note that  $\mathcal{F}_\gamma \preceq (L\mathcal{Q})_\gamma = (\mathcal{Q} \circ \varphi^{-1})_\gamma$  since  $\mathcal{F}$  is a set of *L-upper smooth* functions. Observe that, for all  $\mathcal{F}$ ,

$$\begin{aligned} \mathcal{E}(\mathcal{F} \circ \varphi) &= \left\{ w \mid w [e^{-f \circ \varphi - m(f \circ \varphi)}] \geq 1, \forall f \in \mathcal{F} \right\} \\ &= \left\{ w \mid w \circ \varphi^{-1} [e^{-f - m(f)}] \geq 1, \forall f \in \mathcal{F} \right\} \\ &= \left\{ \tilde{w} \circ \varphi \mid \tilde{w} [e^{-f - m(f)}] \geq 1, \forall f \in \mathcal{F} \right\} \\ &= \mathcal{E}(\mathcal{F}) \circ \varphi, \end{aligned}$$

where  $w$  and  $\tilde{w}$  range over  $\mathcal{M}_+(\mathbb{R}^d)$ . Thus, by Lemma 3, we have  $\mathcal{E}(\mathcal{F}_\gamma) \supset \mathcal{E}((\mathcal{Q} \circ \varphi^{-1})_\gamma) = \mathcal{E}(\mathcal{Q}_{\gamma \circ \varphi}) \circ \varphi^{-1}$ . This proves the inclusion. Now we also have

$$\begin{aligned} C(\gamma, \mathcal{F}_\gamma) &= \inf_{w \in \mathcal{E}(\mathcal{F}_\gamma)} \ln w [e^{-\gamma}] \\ &\leq \inf_{w \in \mathcal{E}(\mathcal{Q}_{\gamma \circ \varphi}) \circ \varphi^{-1}} \ln w [e^{-\gamma}] \\ &= \inf_{w \in \mathcal{E}(\mathcal{Q}_{\gamma \circ \varphi})} \ln w \circ \varphi^{-1} [e^{-\gamma}] \\ &= \inf_{w \in \mathcal{E}(\mathcal{Q}_{\gamma \circ \varphi})} \ln w [e^{-\gamma \circ \varphi}] \\ &= C(\gamma \circ \varphi, \mathcal{Q}_{\gamma \circ \varphi}), \end{aligned}$$

which yields the inequality. ■

This lemma shows that, as long as we consider the envelopes of upper smooth functions  $\mathcal{F}$ , it suffices for bounding them to evaluate the envelopes of penalized quadratic functions  $\mathcal{Q}_\gamma$ .

Further, according to the lemma below, we can restrict ourselves to one-dimensional parametric models

w.l.o.g. if the penalty functions  $\gamma$  is separable. Here,  $\gamma$  is said to be separable if and only if it can be written in the form of  $\gamma(\theta) = \sum_{j=1}^d \gamma_j(\theta_j)$ .

**Lemma 5 (Separability)** *Suppose that  $\gamma$  is separable. Then, the envelope complexity of  $\mathcal{Q}_\gamma$  is bounded by a separable function, i.e.,*

$$C(\gamma, \mathcal{Q}_\gamma) \leq \sum_{j=1}^d C(\gamma_j, \mathcal{Q}_{\gamma_j}^1),$$

where  $\mathcal{Q}^1$  is the set of normalized one-dimensional quadratic functions with curvature one,  $\mathcal{Q}^1 = \{x \in \mathbb{R} \mapsto \frac{1}{2}(x - u)^2 \mid u \in \mathbb{R}\}$ .

**Proof** Note that all  $f \in \mathcal{Q}_\gamma$  is separable, i.e.,  $f(\theta) = \sum_{j=1}^d f_j(\theta_j)$  where  $f_j \in \mathcal{Q}_{\gamma_j}^1$  and  $\gamma(\theta) = \sum_{j=1}^d \gamma_j(\theta_j)$ . Let  $\mathcal{E}^d = \mathcal{E}(\mathcal{Q}_{\gamma_1}^1) \otimes \cdots \otimes \mathcal{E}(\mathcal{Q}_{\gamma_d}^1)$ . Then we have

$$\begin{aligned} C(\gamma, \mathcal{Q}_\gamma) &= \inf_{w \in \mathcal{E}(\mathcal{Q}_\gamma)} \ln w[e^{-\gamma}] \\ &\leq \inf_{w \in \mathcal{E}^d} \ln w[e^{-\gamma}] && \mathcal{E}^d \subset \mathcal{E}(\mathcal{Q}_\gamma) \\ &= \sum_{j=1}^d \inf_{w_j \in \mathcal{E}(\mathcal{Q}_{\gamma_j}^1)} \ln w_j[e^{-\gamma_j}] \\ &= \sum_{j=1}^d C(\gamma_j, \mathcal{Q}_{\gamma_j}^1). \end{aligned}$$

■

**Summary** We have defined the Bayesian envelope and envelope complexity. The envelope complexity  $C(\gamma, \mathcal{F})$  is equal to the Bayesian minimax regret if  $\mathcal{F}$  is the set of penalized logarithmic loss functions. We have shown that the worst-case regrets associated with any  $w$  in the Bayesian envelope  $\mathcal{E}(\mathcal{F})$  are handily bounded and that the Shtarkov complexity is bounded accordingly through the envelope complexity. Most importantly, the envelope complexity satisfies some useful properties such as monotonicity, parametrization invariance and separability. Specifically, the monotonicity differentiates the envelope complexity from the Shtarkov complexity.

## 4 The Spike-and-Tails Prior for High-Dimensional Prediction

We leverage the results of Section 3 to give a Bayesian predictor closely achieving  $\text{LREG}^*(\gamma)$  where  $\gamma(\theta) = \lambda \|\theta\|_1$ , namely, the spike-and-tails prior. Moreover, the resulting predictor also approximately achieves minimax regret without luckiness  $\text{REG}^*(\{\theta : \gamma(\theta) \leq B\})$  where both  $n$  and  $d$  tend to infinity.

### 4.1 Envelope Complexity for $\ell_1$ -Penalties

Let  $\gamma$  be the weighted  $\ell_1$ -norm given by

$$\gamma(\theta) = \lambda \|\theta\|_1, \quad (8)$$

where  $\lambda > 0$ . Let  $\pi_\lambda$  be the spike-and-tails (ST) prior over  $\mathbb{R}^d$  given by

$$\pi_\lambda^{\text{ST}}(d\theta) \propto e^{-\lambda \|\theta\|_1} \prod_{j=1}^d w_\lambda^{\text{ST}}(d\theta_j), \quad (9)$$

$$w_\lambda^{\text{ST}}(dx) = \delta_0(dx) + \frac{e^{\lambda^2/2}}{\lambda^2 e} \mathbf{1}\{|x| \geq \lambda\} dx, \quad (10)$$

where  $\delta_t$  denotes Kronecker's delta measure at  $t$ . We call it the spike-and-tails prior because it consists of a delta measure (spike) and two exponential distributions (tails) as shown in Figure 1.

Then, envelope complexities for quadratic loss functions can be bounded as follows.




Figure 1: Density of the spike-and-tails (ST) prior

**Lemma 6 (Sharp bound for  $\ell_1$ -penalties)** *Take  $\gamma$  as given by (8). Then, we have  $w_\lambda^{\text{ST}} \in \mathcal{E}(\mathcal{Q}_\gamma)$  and*

$$\begin{aligned} d \ln \left( 1 + \frac{e^{-\lambda^2/2}}{\lambda^3(c + o(1))} \right) &\leq C(\gamma, \mathcal{Q}_\gamma) \leq \ln w_\lambda^{\text{ST}} [e^{-\gamma}] \\ &= d \ln \left( 1 + \frac{2e^{-\lambda^2/2}}{\lambda^2 e} \right) \end{aligned}$$

for some constant  $c$ , where  $o(1) \rightarrow 0$  as  $\lambda \rightarrow \infty$ .

**Proof** Consider the logarithmic loss functions of the  $d$ -dimensional standard normal location model, given by  $f_X(\theta) = \frac{1}{2} \|X - \theta\|_2^2 + \frac{d}{2} \ln 2\pi$ ,  $X \in \mathcal{X} = \mathbb{R}^d$  and let  $\mathcal{F} = \{f_X \mid X \in \mathbb{R}^d\}$ . Note that  $\mathcal{F} \preceq \mathcal{Q}$ . Then, the lower bound follows from Lemma 8 in Section A, which is based on Komatu (1955). with  $S(\gamma) \leq C(\gamma, \mathcal{F}_\gamma) \leq C(\gamma, \mathcal{Q}_\gamma)$ .

Since  $\gamma$  is separable, Lemma 4 allows us to restrict ourselves to the case of  $d = 1$ . Let  $f_u(\theta) = \frac{1}{2} (\theta - u)^2 + \lambda |\theta|$  ( $u \in \mathbb{R}$ ) be the penalized loss functions in  $\mathcal{Q}_\gamma$ . Note that  $m(f_u) = \frac{1}{2} u^2$  if  $|u| \leq \lambda$ , and  $m(f_u) = \lambda |u| - \frac{1}{2} \lambda^2$  otherwise.

Let  $c$  and  $t$  be positive real numbers. Let  $w = \delta + cU$  be a measure over the real line, where  $\delta$  denotes the delta measure and  $U$  denotes the Lebesgue measures restricted to  $[-\lambda, \lambda]^c = \mathbb{R} \setminus [-t, t]$ . That is, we have  $w(E) = \mathbb{1}_{0 \in E} + c|E \setminus [-t, t]|$  for measurable sets  $E \subset \mathbb{R}$ . Then we have

$$\ln w[e^{-\gamma}] = \ln \left( 1 + \frac{2c}{\lambda} e^{-t\lambda} \right). \quad (11)$$

We want to minimize (11) with respect to  $w \in \mathcal{E}(\mathcal{Q}_\gamma)$ . Recall that, to have  $w \in \mathcal{E}(\mathcal{Q}_\gamma)$ , it suffices for  $c$  and  $t$  to satisfy  $w[e^{-fu}] \geq e^{-m(fu)}$  for all  $u \in \mathbb{R}$ . Below, we only care about the case  $u \geq \lambda$  since it is symmetric with respect to  $u$  and trivially we have  $w[e^{-fu}] \geq \delta[e^{-fu}] \geq e^{-m(fu)}$  for all  $u \in [-\lambda, \lambda]$ . Thus, we assume  $x \stackrel{\text{def}}{=} u - \lambda \geq 0$ .

Now, observe that

$$\begin{aligned} w[e^{-fu}] &= e^{-\frac{1}{2}u^2} + ce^{-t\lambda} \left( \int_{-\infty}^{-t} + \int_t^{\infty} \right) e^{-\frac{1}{2}(\theta-u)^2} d\theta \\ &\geq e^{-\frac{1}{2}u^2} + ce^{-t\lambda} \int_t^{\infty} e^{-\frac{1}{2}(\theta-u)^2} d\theta \\ &= e^{-m(fu)} \left( e^{-\frac{1}{2}x^2} + c \int_{t-x}^{\infty} e^{-\frac{1}{2}y^2} dy \right) \\ &\stackrel{\text{def}}{=} A(x)e^{-m(fu)}. \end{aligned}$$

Thus, a sufficient condition for  $w \in \mathcal{E}(\mathcal{Q}_\gamma)$  is that  $A'(x) = ce^{-\frac{1}{2}(t-x)^2} - xe^{-\frac{1}{2}x^2} \geq 0$ , which is satisfied with  $c = \frac{1}{t} \exp(\frac{1}{2}t^2 - 1)$ . Evaluating (11) at  $t = \lambda$  yields the ST pre-prior  $w = w_\lambda^{\text{ST}}$ .

Therefore, we have  $w_\lambda^{\text{ST}} \in \mathcal{E}(\mathcal{Q}_\gamma)$  and the upper bound is shown. The equality is a result of straightforward calculation of  $\ln w[e^{-\gamma}]$ .  $\blacksquare$

According to Lemma 6, the ST prior bounds the envelope complexity in a quadratic rate as  $\lambda \rightarrow \infty$ . The exponent,  $-\frac{1}{2}\lambda^2/2$ , is optimally sharp since the lower bound  $C(\gamma, \mathcal{Q}_\gamma) = \Omega(d \exp[-\frac{1}{2}\lambda^2]/\lambda^3)$  has the same exponent.

Finally, we give similar inequalities for general smooth loss functions based on Lemma 4. Let  $\pi_{\lambda, L}^{\text{ST}}$  and  $w_{\lambda, L}^{\text{ST}}$  be the scale-corrected ST (pre-) prior given by

$$\begin{aligned} \pi_{\lambda, L}^{\text{ST}}(d\theta) &= \pi_{\lambda/\sqrt{L}}^{\text{ST}}(\sqrt{L}d\theta), \\ w_{\lambda, L}^{\text{ST}}(d\theta) &= w_{\lambda/\sqrt{L}}^{\text{ST}}(\sqrt{L}d\theta). \end{aligned}$$

The following is a direct corollary of Lemma 4, 5, 6 and 3.

**Corollary 1** *If all  $f \in \mathcal{F}$  is  $L$ -upper smooth with respect to  $\theta$ , and if  $\gamma$  is given by (8), then  $w_{\lambda, L}^{\text{ST}} \in \mathcal{E}(\mathcal{F}_\gamma)$  and therefore*

$$C(\gamma, \mathcal{F}_\gamma) \leq \ln w_{\lambda, L}^{\text{ST}}[e^{-\gamma}] = d \ln \left( 1 + \frac{2L}{e\lambda^2} e^{-\frac{1}{2L}\lambda^2} \right).$$

## 4.2 Regret Bound with the ST Prior

Now, we utilize Corollary 1 for bounding actual prediction performance of the ST prior. Below, we consider the scenario of the online-learning under  $\ell_1$ -constraint.

**Setup** Let  $X^n = (X_1, \dots, X_n) \in \mathcal{X}^n$  be a sequence of outcomes. Let  $f_X$  be a logarithmic loss function such that  $\int e^{-f_X(\theta)} d\nu(X) \leq 1$ . Then, the conditional Bayesian pre-posterior with respect to  $w \in \mathcal{M}_+(\mathbb{R}^d)$  given  $X^t$  ( $0 \leq t \leq n$ ) is given by

$$w(d\theta|X^t) = w(d\theta) \prod_{i=1}^t \exp\{-f_{X_i}(\theta)\}.$$

The online worst-case regret of the predictor is defined as

$$\text{REG}_n(w|\mathcal{H}) \stackrel{\text{def}}{=} \sup_{X^n \in \mathcal{X}^n, \theta^* \in \mathcal{H}} \sum_{t=1}^n \{f_{X_t}(w(\cdot|X^{t-1})) - f_{X_t}(\theta^*)\}. \quad (12)$$

The following theorem shows that the online regret of the ST prior is adaptively<sup>1</sup> minimax.

### Theorem 7 (Minimaxity of the ST priors)

*Suppose that  $f_{X_i}$  are  $L$ -upper smooth and logarithmic. Let  $\mathcal{H}_B = \{\theta \in \mathbb{R}^d \mid \|\theta\|_1 \leq B\}$  and take  $\lambda = \sqrt{2Ln \ln(d/\sqrt{Ln})}$ . Then, with  $w(1) = \ln(d/\sqrt{n}) = o(n)$ , we have*

$$\text{REG}_n(w_{\lambda, Ln}^{\text{ST}}|\mathcal{H}_B) \leq B \sqrt{2Ln \ln \frac{d}{\sqrt{Ln}}} (1 + o(1))$$

*for all  $B > 0$ . Moreover, this is adaptive minimax rate and not improvable more than a factor of two even if  $B$  is fixed and non-Bayesian predictors are involved.*

**Proof** Let  $f_{X^n}$  be the cumulative loss,  $f_{X^n} = \sum_{i=1}^n f_{X_i}$ , and observe that  $f_{X^n}$  is  $Ln$ -upper smooth and logarithmic. Let  $\mathcal{F} = \{f_{X^n} \mid X^n \in \mathcal{X}^n\}$  and  $\gamma(\theta) = \lambda \|\theta\|_1$ . Also, let  $\gamma_0$  be the indicator penalty of the set  $\mathcal{H}_B$  such that  $\gamma_0(\theta) = 0$  if and only if  $\theta \in \mathcal{H}_B$  and otherwise  $\gamma_0(\theta) = \infty$ . Then, we have  $\text{REG}_n(w|\mathcal{H}_B) = \text{LREG}(w|\gamma_0)$  where LREG is taken with respect to  $f_{X^n}$ . Now, observe that

$$\begin{aligned} \text{LREG}(w_{\lambda, Ln}^{\text{ST}}|\gamma_0) &\leq \text{LREG}(w_{\lambda, Ln}^{\text{ST}}|\gamma - \lambda B) \\ &\quad (\because \gamma_0 \geq \gamma - \lambda B) \\ &\leq \ln w_{\lambda, Ln}^{\text{ST}}[e^{-\gamma + \lambda B}], \\ &\quad (\because \text{Theorem 1}) \\ &= \lambda B + \ln w_{\lambda, Ln}^{\text{ST}}[e^{-\gamma}], \end{aligned}$$

<sup>1</sup>The term ‘adaptive’ here refers to the fact that the ST prior with a fixed parameter  $\lambda$  achieves the minimax rate for any radius  $B$  of the  $\ell_1$ -ball.

which, combined with Corollary 1 where  $\lambda = \sqrt{2Ln \ln(d/\sqrt{Ln})}$ , yields the asymptotic equality. The proof of the lower bound is adopted from the existing analysis on the minimax *risk* (see Section B for the rigorous proof and Section 6.5 for detailed discussions). ■

## 5 Visual Comparison of the ST Prior and the Tilted Jeffreys Prior

Now, we verify the results on the  $\ell_1$ -regularization obtained above. In particular, we compare the worst-case regrets achievable with Bayesian predictors to the minimax regret, i.e., the Shtarkov complexity.

**Setting** We adopted the one-dimensional quadratic loss functions with curvature one,  $q \in \mathcal{Q}^1$ , and the  $\ell_1$ -penalty function,  $\gamma(\theta) = \lambda|\theta|$ . We varied the penalty weight  $\lambda$  from  $10^{-1}$  to  $10^1$  and observed how the worst-case regret of each Bayesian predictor changes. Specifically, we employed the spike-and-tails (ST) prior (10) and the tilted Jeffreys prior for the predictors. Note that, in this case, the tilted Jeffreys prior is nothing more than the double exponential prior given by  $\pi_\lambda^{\text{Jeff}}(d\theta) = \frac{\lambda}{2} e^{-\lambda|\theta|} d\theta$ .

**Results** In Figure 2, the worst-case regrets of the ST prior and the tilted Jeffreys prior are shown along with the minimax regret (Optimal). While the regret of the tilted Jeffreys prior is almost same as the optimal regret where  $\lambda$  is small, it performs poorly where  $\lambda$  is large. On the other hand, the ST prior performs robustly well in the entire range of  $\lambda$ . Specifically, it converges to zero quadratically where  $\lambda$  is large. Therefore, since one must take  $\lambda$  sufficiently large if  $d$  is large, it is implied that the ST prior is a better choice than the tilted Jeffreys prior in high-dimensional settings.

## 6 Implications and Discussions

In this section, we discuss interpretations of the results and present solutions to some technical difficulties.

### 6.1 Gap between $\text{LREG}^*$ and $\text{LREG}^{\text{Bayes}}$

One may wonder if there exists a prior that achieves the lower bound  $\text{LREG}^*(\gamma)$  where  $\gamma(\theta) = \lambda \|\theta\|_1$ ,  $\lambda > 0$ . Unfortunately, the answer is negative. With a similar technique of higher-order differentiations used by Hedayati and Bartlett (2012), we can show that, if  $\gamma$  is convex and not differentiable like the  $\ell_1$ -norm, then

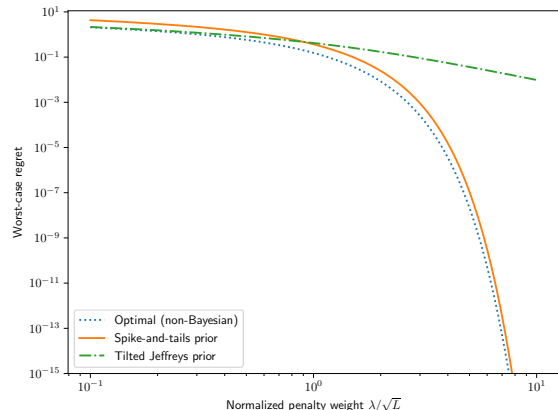


Figure 2: Worst-case regrets of the spike-and-tails (ST) prior and the tilted Jeffreys prior

the gap is nonzero, i.e.,  $\text{LREG}^*(\gamma) < \text{LREG}^{\text{Bayes}}(\gamma)$ . The detailed statement and the proof is in Section C.

### 6.2 Infinite-dimensional Models

If the dimensionality  $d$  of the parameter space is countably infinite, the minimax regret  $\text{REG}^*(\mathcal{H}_B)$  with any nonzero radius  $B$  diverges. In this case, one may apply different penalty weights to different dimensions. For instance, taking the penalty as  $\gamma(\theta) = \sum_{j=1} \lambda_j |\theta_j|$  for  $\lambda_j = \sqrt{2L \text{Ln}\{j \text{Ln} j\}}$  and  $\text{Ln} x = \ln \max\{e, x\}$ , the separability of the envelope complexity guarantees that  $C(\gamma, \mathcal{F}_\gamma) \leq \sum_{j=1}^{\infty} (j \text{Ln}^2 j)^{-1} < +\infty$ . Then, the corresponding countably-infinite tensor product of the one-dimensional ST prior  $\pi_{\{\lambda_j\}}^{\text{ST}}(d\theta) = \prod_{j=1}^{\infty} \pi_{\lambda_j}^{\text{ST}}(d\theta_j)$  gives a finite regret with respect to the infinite-dimensional models  $\mathcal{H} = \{\theta \in \mathbb{R}^{\mathbb{N}} \mid \gamma(\theta) \leq B\}$ .

### 6.3 Comparison to the Titled Jeffreys Priors and Others

There have been previous studies on the minimax regret with Bayesian predictors (Takeuchi and Barron, 1998, 2013; Watanabe and Roos, 2015; Xie and Barron, 2000). In these studies, the Bayesian predictor based on the Jeffreys prior (namely Jeffreys predictor) is proved to attain minimax-regret asymptotically under some regularity conditions. The tilted Jeffreys prior, which takes the effect of penalization  $\gamma$  into consideration, is given by Grünwald (2007) as  $\pi_{\text{Jeff}}(d\theta) \propto d\theta \sqrt{\det I(\theta)} e^{-\gamma(\theta)}$ , where  $I(\theta)$  denotes the Fisher information matrix. In the case of quadratic loss functions  $\mathcal{Q}$ , as the Fisher information is equal to identity, we have  $\pi_{\text{Jeff}}(d\theta) \propto e^{-\gamma(\theta)} d\theta$ . Therefore, it implies that taking the uniform pre-prior  $w(d\theta) \propto d\theta$  is good for smooth models under the conventional large-sample limit. This is in very strong contrast with

our result, where completely nonuniform pre-prior  $w_\lambda^{\text{ST}}$  performs better with high-dimensional models.

#### 6.4 Comparison to Online Convex Optimization

So far, we have considered the luckiness minimax regret, which leads to the adaptive minimax regret. Perhaps surprisingly, our minimax regret bound coincides with the results given in the literature of online convex optimization, where different assumptions on the loss functions and predictors are made. Specifically, with  $\lambda = \sqrt{2L \ln d}$ , the regret bound is reduced to  $\sqrt{2L \ln d} + 1/e$ . This coincides with the standard no-regret rates of online learning such as Hedge algorithm (Freund and Schapire, 1997), online mirror descent (Nemirovski et al., 2009) and high-dimensional online regression (Gerchinovitz and Yu, 2014), where  $L$  is referred to as the number of trials  $T$  and  $d$  is referred to as the number of experts or dimensions  $n$ . Moreover, with  $\lambda = 1$ , the regret bound is reduced to  $O(d \ln L)$ . This is equal to the minimax-regret rate achieved under large-sample asymptotics such as in Hazan et al. (2007); Cover (2011).

Note that, the conditions assumed in those two regimes are somewhat different. In our setting, loss functions are assumed to be upper smooth and satisfy some normalizing condition to be logarithmic losses, while the boundedness and convexity of loss functions is often assumed in online learning. Moreover, we have employed Bayesian predictors, whereas more simple online predictors are typically used in the context of the online learning.

#### 6.5 Comparison to Minimax Risk over $\ell_1$ -balls

In the literature of high-dimensional statistics, the minimax rate of *statistical risk* is also achieved with  $\ell_1$ -regularization (Donoho and Johnstone, 1994), when the true parameter  $\theta$  is in the unit  $\ell_1$ -ball. Although both risk and regret are performance measures of prediction, there are two notable difference. One is that risks are calculated under some assumptions on true statistical distribution, whereas regrets are defined without any assumptions on data. The other is that risks are typically considered with in-model predictor, i.e., predictors are restricted to a given model, whereas regrets are often considered with out-model predictors such as Bayesian predictors and online predictors. Therefore, the minimax regret can be regarded as a more agnostic complexity measure than the minimax risk.

If we assume Gaussian noise models and adopt the logarithmic loss functions, the minimax rate of the risk is

given as  $\sqrt{2L \ln d / \sqrt{L}}$  according to Donoho and Johnstone (1994). Interestingly, this is same with the rate of the regret bound given by Theorem 7 where  $L = Ln$ . Moreover, the minimax-risk optimal penalty weights  $\lambda$  is also minimax-regret optimal in this case. Therefore, if the dimensionality  $d$  is large enough compared to  $L$  ( $n$  in case of online-learning), making no distributional assumption on data costs nothing in terms of the minimax rate.

## 7 Conclusion

In this study, we presented a novel characterization of the minimax regret for logarithmic loss functions, called the envelope complexity, with  $\ell_1$ -regularization problems. The virtue of the envelope complexity is that it is much easier to evaluate than the minimax regret itself and able to produce upper bounds systematically. Then, using the envelope complexity, we have proposed the spike-and-tails (ST) prior, which almost achieves the luckiness minimax regret against smooth loss functions under  $\ell_1$ -penalization. We also show that the ST prior actually adaptively achieves the 2-approximate minimax regret under high-dimensional asymptotics  $\omega(1) = \ln d / \sqrt{n} = o(n)$ . In the experiment, we have confirmed our theoretical results: The ST prior outperforms the tilted Jeffreys prior where the dimensionality  $d$  is high, whereas the tilted Jeffreys prior is optimal if  $n \gg d$ .

**Limitation and future work** The present work is relying on the assumption of the smoothness and logarithmic property on the loss functions. The smoothness assumption may be removed by considering the smoothing effect of stochastic algorithms like stochastic gradient descent as in Kleinberg et al. (2018). As for the logarithmic assumption, it will be generalized to non-logarithmic loss functions with the help of tools that have been developed in the literature of information theory such as in Yamanishi (1998). Finally, since our regret bound with the ST prior is quite simple (there are only the smoothness  $L$  and the radius  $B$  except with the logarithmic term), applying these results to concrete models such as deep learning models would be interesting future work as well as the comparison to the existing generalization error bounds.

#### Acknowledgements

This work is partially supported by JST, CREST and MEXT KAKENHI 23240019.

#### References

Barron, A., Roos, T., and Watanabe, K. (2014). Bayesian properties of normalized maximum like-



- likelihood and its fast computation. In *IEEE International Symposium on Information Theory - Proceedings*.
- Cesa-Bianchi, N., Conconi, A., and Gentile, C. (2004). On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057.
- Cesa-Bianchi, N. and Gentile, C. (2008). Improved risk tail bounds for on-line algorithms. *IEEE Transactions on Information Theory*, 54(1):386–390.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge university press.
- Cover, T. M. (2011). Universal portfolios. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pages 181–209. World Scientific.
- Donoho, D. L. and Johnstone, I. M. (1994). Minimax risk over  $p$ -balls for  $p$ -error. *Probability Theory and Related Fields*, 99(2):277–303.
- Freund, Y. and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139.
- Gerchinovitz, S. and Yu, J. Y. (2014). Adaptive and optimal online linear regression on 1-balls. *Theoretical Computer Science*, 519:4–28.
- Grünwald, P. D. (2007). *The minimum description length principle*. MIT press.
- Hazan, E., Agarwal, A., and Kale, S. (2007). Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192.
- Hedayati, F. and Bartlett, P. L. (2012). The optimality of jeffreys prior for online density estimation and the asymptotic normality of maximum likelihood estimators. In *Conference on Learning Theory*, pages 7–1.
- Kleinberg, R., Li, Y., and Yuan, Y. (2018). An alternative view: When does sgd escape local minima? *arXiv preprint arXiv:1802.06175*.
- Komatu, Y. (1955). Elementary inequalities for mills ratio. *Rep. Statist. Appl. Res. Un. Jap. Sci. Engrs*, 4:69–70.
- Koolen, W. M., Malek, A., and Bartlett, P. L. (2014). Efficient minimax strategies for square loss games. In *Advances in Neural Information Processing Systems*, pages 3230–3238.
- Littlestone, N. (1989). From on-line to batch learning. In *Proceedings of the second annual workshop on Computational learning theory*, pages 269–284.
- Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. (2009). Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609.
- Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, 14(5):465–471.
- Rissanen, J. J. (1996). Fisher information and stochastic complexity. *IEEE transactions on information theory*, 42(1):40–47.
- Shtar’kov, Y. M. (1987). Universal sequential coding of single messages. *Problemy Peredachi Informatsii*, 23(3):3–17.
- Takeuchi, J. and Barron, A. R. (1998). Asymptotically minimax regret by bayes mixtures. In *IEEE International Symposium on Information Theory - Proceedings*.
- Takeuchi, J. and Barron, A. R. (2013). Asymptotically minimax regret by bayes mixtures for non-exponential families. In *Information Theory Workshop (ITW), 2013 IEEE*, pages 1–5. IEEE.
- Watanabe, K. and Roos, T. (2015). Achievability of asymptotic minimax regret by horizon-dependent and horizon-independent strategies. *The Journal of Machine Learning Research*, 16(1):2357–2375.
- Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688.
- Xie, Q. and Barron, A. R. (2000). Asymptotic minimax regret for data compression, gambling, and prediction. *IEEE Transactions on Information Theory*, 46(2):431–445.
- Yamanishi, K. (1998). A decision-theoretic extension of stochastic complexity and its applications to learning. *IEEE Transactions on Information Theory*, 44(4):1424–1439.