

# Uncertainty-Aware Bayes' Rule and Its Applications

Shixiong Wang

**Abstract**—Bayes' rule has enabled innumerable powerful algorithms of statistical signal processing and statistical machine learning. However, when there exist model misspecifications in prior distributions and/or data distributions, the direct application of Bayes' rule is questionable. Philosophically, the key is to balance the relative importance of prior and data distributions when calculating posterior distributions: if prior (resp. data) distributions are overly conservative, we should upweight the prior belief (resp. data evidence); if prior (resp. data) distributions are overly opportunistic, we should downweight the prior belief (resp. data evidence). This paper derives a generalized Bayes' rule, called uncertainty-aware Bayes' rule, to technically realize the above philosophy, i.e., to combat the model uncertainties in prior distributions and/or data distributions. Simulated and real-world experiments on classification and estimation showcase the superiority of the presented uncertainty-aware Bayes' rule over the conventional Bayes' rule: In particular, the uncertainty-aware Bayes classifier, the uncertainty-aware Kalman filter, the uncertainty-aware particle filter, and the uncertainty-aware interactive-multiple-model filter are suggested and validated.

**Index Terms**—Bayes' Rule, Uncertainty Awareness, Entropy Method, Signal Processing, Machine Learning.

## I. INTRODUCTION

### A. Background

Bayes' rule (or Bayes' theorem) has countless successful applications in statistical signal processing and statistical machine learning such as Kalman filter, particle filter, Bayesian hypothesis testing, naive Bayes classifier, Bayesian optimization, Bayesian bandits, and Bayesian networks. Consider a data-generating distribution  $p_{\theta_0}(\mathbf{y})$  parametrized by unknown  $\theta_0 \in \Theta \subseteq \mathbb{R}^d$  where  $\Theta$  is the parameter space. We aim to estimate the true value  $\theta_0$  using  $n$  collected samples  $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\} \subset \mathcal{Y}$  from  $p_{\theta_0}(\mathbf{y})$  where  $\mathcal{Y} \subseteq \mathbb{R}^m$  is the domain of  $p_{\theta_0}(\mathbf{y})$ . Bayes' rule

$$p(\theta|\mathbf{y}) \propto p(\theta) \cdot p(\mathbf{y}|\theta), \quad \forall \mathbf{y} \in \mathcal{Y}, \forall \theta \in \Theta \quad (1)$$

suggests the law of calculating the posterior distribution  $p(\theta|\mathbf{y})$  after observing the sample  $\mathbf{y}$  using the likelihood function  $\theta \mapsto p(\mathbf{y}|\theta) := p_{\theta}(\mathbf{y})$  and the prior distribution  $p(\theta)$ . Philosophically, the prior distribution  $p(\theta)$  encodes our prior belief that  $\theta \in \Theta$  is the true value and the likelihood function  $p(\mathbf{y}|\theta)$  indicates our  $\mathbf{y}$ -data evidence that  $\theta$  is the true value; therefore, the posterior distribution  $p(\theta|\mathbf{y})$  denotes our integrated posterior belief that  $\theta$  is the true value. For a quick review of Bayes' rule, see, e.g., [1].

### B. Problem Statement and Literature Review

In using the conventional Bayes' rule (1), the fundamental assumption is that the data distribution  $\mathbf{y} \mapsto p(\mathbf{y}|\theta)$  for every

$\theta \in \Theta$  and the prior distribution  $p(\theta)$  are accurate: that is, the true value  $\theta_0$  is indeed a realization of  $p(\theta)$  and the data  $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$  are indeed sampled from  $p_{\theta_0}(\mathbf{y})$ . However, this assumption is highly untenable in practice, and as a result, the performance of algorithms relying on the conventional Bayes' rule (1) significantly degrades: For illustrations and justifications from the general Bayesian statistics community, see, e.g., [2]–[6]; for those from the Bayesian statistical signal processing community, see, e.g., [7]–[10].

Facing the model uncertainties before applying Bayes' rule, one way is to improve the modeling accuracy and reduce such uncertainties, for example, by replacing Gaussian distributions with Student's  $t$  distributions when outliers exist [5, Chap. 17], and the other is to withstand the model uncertainties with robust methods, for example, using noninformative priors [2], using robust data distributions [11, Chaps. 4 and 5], [12], [13], or modifying the Bayes' rule itself (e.g.,  $\alpha$ -posterior) [14, Sec. 6.8.5 and Sec. 8.], [6], [15]. While it is not always practically easy to improve the modeling accuracy (because additional information is required, e.g., the knowledge that Student's  $t$  distributions can model measurement outliers), robust methods are attractive to practitioners. According to [11, Sec. 15.2],

the key to the robust design is to balance the relative importance between the prior distribution (i.e., prior belief) and the data distribution (i.e., likelihood function, data evidence).

To clarify further, if the prior distribution is uncertain (i.e., unreliable), we modify (i.e., reduce or increase) the contribution of the prior distribution to the posterior distribution; likewise, if the data distribution is uncertain, we modify the contribution of the data distribution to the posterior distribution. This idea has been empirically/theoretically validated and technically implemented in, e.g., [16], where the maximum entropy scheme is employed to diminish the importance of prior distributions or data distributions, and in, e.g., [14, Sec. 6.8.5 and Sec. 8.], where the  $\alpha$ -posterior scheme<sup>1</sup> is employed to manage the importance of the data distributions using an  $\alpha$ -power operator (the importance of the data distributions is reduced if  $0 < \alpha < 1$  and remains unchanged if  $\alpha = 1$ ).

However, the maximum entropy scheme is computationally complex: we prefer a computationally simpler scheme that has almost the same computational burden as the conventional Bayes' rule and the  $\alpha$ -posterior scheme. To clarify further, we prefer a Bayes-rule-like formula that can function as the maximum entropy scheme. In addition, both the maximum entropy scheme and the  $\alpha$ -posterior scheme cannot technically augment the importance of prior distributions and data

S. Wang is with the Institute of Data Science, National University of Singapore, Singapore 117602 (E-mail: s.wang@u.nus.edu).

<sup>1</sup>The  $\alpha$ -posterior is defined as  $p_{\alpha}(\theta|\mathbf{y}) \propto p(\theta) \cdot p^{\alpha}(\mathbf{y}|\theta)$ ,  $\alpha > 0$ ; cf. (1).

distributions: they can only diminish the importance of prior distributions and/or data distributions because as robust (i.e., conservative) methods, they tend to downweight the information (from prior belief or data evidence) at hands. Nevertheless, when the prior belief (resp. data evidence) is obtained in a conservative manner, it is beneficial if we can upweight the information from the prior belief (resp. data evidence).

In summary, a generalized Bayes' rule that can balance the relative importance between prior belief and data evidence is expected. Furthermore, the new Bayes' rule can not only downweight prior belief and data evidence but also upweight them when required, in a computationally efficient manner.

### C. Contributions

This paper mathematically formalizes and generalizes the philosophy in [11, Sec. 15.2], and an uncertainty-aware Bayes' rule is derived. The generalization roots in moving to an uncertainty-aware design from the robust design in [11, Sec. 15.2]. Uncertainty awareness is more general than robustness because when the information (i.e., prior belief or data evidence) is believed to be opportunistic, we need to pursue robustness and downweight the information. In contrast, when the information is believed to be conservative, we need to abandon robustness and upweight the information.

The contributions of this paper can be itemized as follows.

- 1) We propose a new interpretation method of the conventional Bayes' rule (1); see Lemma 1.
- 2) We propose an uncertainty-aware Bayes' rule  $p_g(\theta|\mathbf{y}) \propto p^\beta(\theta) \cdot p^\alpha(\mathbf{y}|\theta)$ , called  $(\alpha, \beta)$ -posterior, where  $\alpha, \beta \in [0, \infty]$ , that generalizes the conventional Bayes' rule (1); see Theorem 1 and Definition 2.
- 3) We show that for some values of  $\alpha$  and  $\beta$ , the uncertainty-aware Bayes' rule can technically reflect the robustness (i.e., conservatism) philosophy in [11, Sec. 15.2], that is, to function as the maximum entropy scheme and the  $\alpha$ -posterior scheme (i.e., downweighting the prior belief and/or data evidence). For other values, the uncertainty-aware Bayes' rule can reflect the opportunism philosophy and upweight the prior belief and/or data evidence. For details, see Theorems 2 and 3, and Insights 1 and 2.
- 4) We show that an  $\alpha$ -scaled (resp.  $\beta$ -scaled) nominal distribution  $p_\theta^\alpha(\mathbf{y})$  [resp.  $p^\beta(\theta)$ ], after normalization, is closer to the true distribution  $p_{\theta_0}(\mathbf{y})$  [resp.  $p_0(\theta)$ ] than the original nominal distribution  $p_\theta(\mathbf{y})$  [resp.  $p(\theta)$ ] for some  $\alpha \geq 0$  (resp.  $\beta \geq 0$ ), which can justify the superiority of the proposed  $(\alpha, \beta)$ -posterior over the conventional Bayes' posterior; see Theorem 4.
- 5) Applications of the  $(\alpha, \beta)$ -posterior in machine learning and signal processing (e.g., classification and estimation) are discussed; see Subsection II-F. In addition, the power of the  $(\alpha, \beta)$ -posterior on both simulated and real-world data sets is empirically validated; see Section III.

### D. Notations

Random and deterministic quantities are denoted by upright and italic symbols (e.g.,  $\mathbf{y}$  and  $\mathbf{y}$ ), respectively. Let  $\mathbb{R}^d$  denote

the  $d$ -dimensional real space. We use  $p_{\mathbf{y}}(\mathbf{y})$  to denote the probability density (resp. mass) function of  $\mathbf{y}$  if  $\mathbf{y}$  is continuous (resp. discrete); when it is clear from the contexts,  $p(\mathbf{y})$  is used as a shorthand for  $p_{\mathbf{y}}(\mathbf{y})$ . Let  $\text{Ent } p(\mathbf{y})$  denote the entropy of the distribution  $p(\mathbf{y})$ : i.e.,  $\text{Ent } p(\mathbf{y}) := \int p(\mathbf{y}) \cdot -\ln p(\mathbf{y}) d\mathbf{y}$ .<sup>2</sup> Let  $\text{KL}[q(\theta) \| p(\theta)]$  define the Kullback–Leibler divergence of  $q(\theta)$  from  $p(\theta)$ . The Gaussian distribution with mean  $\mu$  and covariance  $\Sigma$  is denoted as  $\mathcal{N}(\mu, \Sigma)$  and its density function as  $\mathcal{N}(\cdot; \mu, \Sigma)$ . The running index set induced by integer  $r$  is denoted as  $[r] := \{1, 2, \dots, r\}$ .

## II. MAIN RESULTS

This paper limits the presentation to the elementary probability theory and avoids using measure-theoretic languages. In addition, throughout the paper, we assume that density functions exist (with respect to the Lebesgue measure) and take values on  $[0, \infty)$ . To reduce the presentation length, the main results are only given under probability density functions; for probability mass functions, one may just change integrals to sums. (Simplified proofs for discrete cases are possible since probability masses take values on  $[0, 1]$ .) To reduce notational clutter, we implicitly mean  $\mathbf{y} \in \mathcal{Y}$  and  $\theta \in \Theta$  throughout the paper, unless otherwise stated.

### A. Generalized Bayes' Rule

We begin with the notion of likelihood distribution.

*Definition 1 (Likelihood Distribution):* Let

$$l_{\mathbf{y}}(\theta) := \frac{p(\mathbf{y}|\theta)}{\int p(\mathbf{y}|\theta) d\theta}, \quad \forall \mathbf{y} \quad (2)$$

denote the *likelihood distribution* of  $\theta$  induced by the likelihood function  $\theta \mapsto p(\mathbf{y}|\theta)$  evaluated at the sample  $\mathbf{y}$ . When it is clear from the contexts, we suppress the notational dependence on  $\mathbf{y}$  and use  $l(\theta)$  as a shorthand for  $l_{\mathbf{y}}(\theta)$ .  $\square$

The likelihood distribution  $l(\theta)$  is a  $\mathbf{y}$ -parametric distribution of  $\theta$ . A direct result of Definition 1 is that the posterior distribution  $p(\theta|\mathbf{y})$  given by the conventional Bayes' rule (1) can also be expressed as

$$p(\theta|\mathbf{y}) \propto p(\theta) \cdot l(\theta). \quad (3)$$

The lemma below gives an interpretation of the origin of the Bayes' rule (3).

*Lemma 1 (Conventional Bayes' Rule):* The posterior distribution  $p(\theta|\mathbf{y})$  given by the Bayes' rule (1) [or (3)] solves

$$\min_{q(\theta)} \text{KL}[q(\theta) \| p(\theta)] + \text{KL}[q(\theta) \| l(\theta)] + \text{Ent } q(\theta). \quad (4)$$

*Proof:* See Appendix A.  $\square$

Lemma 1 suggests that the posterior distribution  $p(\theta|\mathbf{y})$  is an entropy-regularized  $(\frac{1}{2}, \frac{1}{2})$ -KL-barycenter of the prior distribution  $p(\theta)$  and the likelihood distribution  $l(\theta)$ ; intuitively,  $p(\theta|\mathbf{y})$  is a minimum-entropy distribution that is simultaneously close to both the prior distribution (i.e., prior belief) and the likelihood distribution (i.e.,  $\mathbf{y}$ -data evidence).

<sup>2</sup>Unless otherwise stated, in this paper, when we mention entropy, we mean the differential entropy for continuous distributions and the Shannon entropy for discrete ones.

We generalize the optimization problem (4) to

$$\min_{q(\theta)} \alpha_1 \text{KL}[q(\theta) \| p(\theta)] + \alpha_2 \text{KL}[q(\theta) \| l(\theta)] + \alpha_3 \text{Ent } q(\theta), \quad (5)$$

where  $\alpha_1, \alpha_2$ , and  $\alpha_3$  are weights. To let  $q(\theta)$  be close to  $p(\theta)$  and  $l(\theta)$ , without loss of practicality, the following assumption is imposed.

*Assumption 1:* We assume that  $0 \leq \alpha_1 < \infty$ ,  $0 \leq \alpha_2 < \infty$ , and  $-\infty < \alpha_3 < \infty$ .  $\square$

Depending on whether to pursue the maximum or minimum entropy of  $q(\theta)$ ,  $\alpha_3$  can take any finite value on the whole real line  $\mathbb{R}$ . We term the distribution solving the generalized problem (5) the *generalized Bayes' posterior distribution*.

*Theorem 1 (Generalized Bayes' Rule):* The generalized posterior distribution  $p_g(\theta|y)$  solving (5) is given by

$$p_g(\theta|y) \propto p^{\frac{\alpha_1}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta) \cdot l^{\frac{\alpha_2}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta), \quad (6)$$

or equivalently,

$$p_g(\theta|y) \propto p^{\frac{\alpha_1}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta) \cdot p^{\frac{\alpha_2}{\alpha_1 + \alpha_2 - \alpha_3}}(y|\theta), \quad (7)$$

when  $\alpha_3 < \alpha_1 + \alpha_2$ , provided that the right-hand-side terms are integrable on  $\Theta$ . When  $\alpha_3 = \alpha_1 + \alpha_2$ ,  $p_g(\theta|y)$  is an arbitrary distribution supported on the set  $\Theta^*$  where

$$\Theta^* := \underset{\theta}{\operatorname{argmax}} \alpha_1 \ln p(\theta) + \alpha_2 \ln l(\theta) \quad (8)$$

contains all weighted maximum *a-posteriori* estimates.

*Proof:* See Appendix B.  $\square$

In Theorem 1 and throughout this paper, we only consider the case where  $\alpha_3 \leq \alpha_1 + \alpha_2$  for technical simplicity; see the proof of Theorem 1 for details. Eq. (8) can be understood as an  $(\alpha_1, \alpha_2)$ -weighted maximum posterior estimation. Motivated by the generalized posterior distribution in (6), we give the definition of the  $(\alpha, \beta)$ -posterior.

*Definition 2 (( $\alpha, \beta$ )-Posterior):* The  $(\alpha, \beta)$ -posterior induced by the prior distribution  $p(\theta)$  and the likelihood distribution  $l(\theta)$  is defined as

$$p_g(\theta|y) \propto p^\beta(\theta) \cdot l^\alpha(\theta), \quad (9)$$

where  $0 \leq \alpha, \beta < \infty$ , if the right-hand-side term is integrable on  $\Theta$ . When  $\alpha$  and  $\beta$  are infinity, the  $(\alpha, \beta)$ -posterior  $p_g(\theta|y)$  is defined as an arbitrary distribution supported on  $\Theta^*$  where  $\Theta^* := \operatorname{argmax}_\theta \alpha_1 \ln p(\theta) + \alpha_2 \ln l(\theta)$ .  $\square$

Note that, compared with (6),  $\beta := \frac{\alpha_1}{\alpha_1 + \alpha_2 - \alpha_3} \in [0, \infty)$  and  $\alpha := \frac{\alpha_2}{\alpha_1 + \alpha_2 - \alpha_3} \in [0, \infty)$ , if  $\alpha_3 < \alpha_1 + \alpha_2$ . When  $\alpha_3 \uparrow (\alpha_1 + \alpha_2)$ ,  $\alpha$  and  $\beta$  simultaneously tend to infinity.

*Remark 1:* If  $\alpha_3 \geq 0$  is additionally required, then  $\alpha + \beta \geq 1$  must be appended in Definition 2. On the other hand, if we allow  $\alpha_1$  and  $\alpha_2$  to be negative, then  $\beta$  and  $\alpha$  can be negative as well. Negative values for  $\alpha_1$  (resp.  $\alpha_2$ ) imply that  $q(\theta)$  is expected to be far away from  $p(\theta)$  [resp.  $l(\theta)$ ]. Without loss of practicality, this paper focuses on the specifications in Assumption 1.  $\square$

Below we give several motivational examples of the generalized Bayes' rule (6) or (9); the first two are well-established in existing literature of Bayesian statistics.

*Example 1 (Conventional Bayes' Posterior):* The conventional Bayes' rule (1) [resp. (3)] is a special case of the generalized Bayes' rule (7) [resp. (6) or (9)] when  $\alpha_1 = \alpha_2 = \alpha_3$  or  $\alpha = \beta = 1$ .  $\square$

*Example 2 ( $\alpha$ -Posterior):* When  $\alpha_2 = \alpha_3$  and  $\alpha_1 > 0$ , (6) reduces to

$$p_g(\theta|y) \propto p(\theta) \cdot l^{\frac{\alpha_2}{\alpha_1}}(\theta). \quad (10)$$

By letting  $\alpha := \frac{\alpha_2}{\alpha_1}$ , we obtain the  $\alpha$ -posterior

$$p_g(\theta|y) \propto p(\theta) \cdot l^\alpha(\theta), \quad (11)$$

where  $0 \leq \alpha < \infty$ . When  $\alpha_1 = 0$ ,  $p_g(\theta|y)$  is an arbitrary distribution supported on  $\Theta^*$  where  $\Theta^* := \operatorname{argmax}_\theta \ln l(\theta)$  contains all maximum likelihood estimates.  $\square$

The  $\alpha$ -posterior in (11) is a well-established proposal in Bayesian statistics; see, e.g., [14, Sec. 6.8.5 and Sec. 8.6], [6], [15]. Given  $0 < \alpha < 1$  and several other technical regularity conditions (e.g.,  $\Theta^*$  is a singleton),  $\alpha$ -posteriors can be shown to have posterior consistency (see [14, Thm. 6.54, Ex. 8.44] and [15]), asymptotic normality (see [6, Thm. 1]), and robustness against likelihood-model misspecifications (see [6, Sec. 4]); however, these conditions might be practically restrictive.

*Example 3 ( $\beta$ -Posterior):* When  $\alpha_1 = \alpha_3$  and  $\alpha_2 > 0$ , (6) reduces to

$$p_g(\theta|y) \propto p^{\frac{\alpha_1}{\alpha_2}}(\theta) \cdot l(\theta). \quad (12)$$

By letting  $\beta := \frac{\alpha_1}{\alpha_2}$ , we obtain the  $\beta$ -posterior<sup>3</sup>

$$p_g(\theta|y) \propto p^\beta(\theta) \cdot l(\theta), \quad (13)$$

where  $0 \leq \beta < \infty$ . When  $\alpha_2 = 0$ ,  $p_g(\theta|y)$  is an arbitrary distribution supported on  $\Theta^*$  where  $\Theta^* := \operatorname{argmax}_\theta \ln p(\theta)$  contains all maximum prior estimates.  $\square$

Compared with the  $\alpha$ -posterior in Example 2 that modifies the likelihood distribution  $l(\theta)$ , the  $\beta$ -posterior modifies the prior distribution  $p(\theta)$ .

*Example 4 ( $\gamma$ -Posterior):* When  $\alpha_1 = \alpha_2$  and  $2\alpha_1 > \alpha_3$ , (6) reduces to

$$p_g(\theta|y) \propto p^{\frac{\alpha_1}{2\alpha_1 - \alpha_3}}(\theta) \cdot l^{\frac{\alpha_1}{2\alpha_1 - \alpha_3}}(\theta), \quad (14)$$

By letting  $\gamma := \frac{\alpha_1}{2\alpha_1 - \alpha_3}$ , we obtain the  $\gamma$ -posterior<sup>4</sup>

$$p_g(\theta|y) \propto p^\gamma(\theta) \cdot l^\gamma(\theta), \quad (15)$$

where  $0 \leq \gamma < \infty$ . When  $2\alpha_1 = \alpha_3$ ,  $p_g(\theta|y)$  is an arbitrary distribution supported on  $\Theta^*$  where  $\Theta^* := \operatorname{argmax}_\theta \ln p(\theta) + \ln l(\theta)$  contains all usual maximum *a-posteriori* estimates.  $\square$

Compared with the  $\alpha$ -posterior and the  $\beta$ -posterior, the  $\gamma$ -posterior modifies both the prior distribution and the likelihood distribution.

*Example 5 ( $\alpha$ -Likelihood):* When  $\alpha_1 = 0$  and  $\alpha_2 > \alpha_3$ , (6) reduces to

$$p_g(\theta|y) \propto l^{\frac{\alpha_2}{\alpha_2 - \alpha_3}}(\theta). \quad (16)$$

By letting  $\alpha := \frac{\alpha_2}{\alpha_2 - \alpha_3}$ , we obtain the  $\alpha$ -likelihood

$$p_g(\theta|y) \propto l^\alpha(\theta), \quad (17)$$

<sup>3</sup>To differentiate with the  $\alpha$ -posterior, we name it using  $\beta$ .

<sup>4</sup>To differentiate with the  $\alpha$ - and  $\beta$ -posterior, we name it using  $\gamma$ .

where  $0 \leq \alpha < \infty$ . When  $\alpha_2 = \alpha_3$ ,  $p_g(\theta|\mathbf{y})$  is an arbitrary distribution supported on  $\Theta^*$  where  $\Theta^* := \operatorname{argmax}_{\theta} \ln l(\theta)$  contains all maximum likelihood estimates.  $\square$

In the generalized Bayes' posterior (17) induced by the  $\alpha$ -likelihood  $l^\alpha(\theta)$ , the prior distribution  $p(\theta)$  is completely ignored. To clarify further, we only trust the likelihood distribution with the confidence level  $\alpha$  and absolutely distrust the prior distribution.

*Example 6 ( $\alpha$ -Prior):* When  $\alpha_2 = 0$  and  $\alpha_1 > \alpha_3$ , (6) reduces to

$$p_g(\theta|\mathbf{y}) \propto p^{\frac{\alpha_1}{\alpha_1 - \alpha_3}}(\theta). \quad (18)$$

By letting  $\alpha := \frac{\alpha_1}{\alpha_1 - \alpha_3}$ , we obtain the  $\alpha$ -prior

$$p_g(\theta|\mathbf{y}) \propto p^\alpha(\theta), \quad (19)$$

where  $0 \leq \alpha < \infty$ . When  $\alpha_1 = \alpha_3$ ,  $p_g(\theta|\mathbf{y})$  is an arbitrary distribution supported on  $\Theta^*$  where  $\Theta^* \in \operatorname{argmax}_{\theta} \ln p(\theta)$  contains all maximum prior estimates.  $\square$

In the generalized Bayes' posterior (19) induced by the  $\alpha$ -prior  $p^\alpha(\theta)$ , the likelihood distribution  $l(\theta)$  is completely ignored. To clarify further, we only trust the prior distribution with the confidence level  $\alpha$  and absolutely distrust the likelihood distribution.

*Example 7 ( $\alpha$ -Pooled Posterior):* When  $\alpha_3 = 0$ , (6) reduces to

$$p_g(\theta|\mathbf{y}) \propto p^{\frac{\alpha_1}{\alpha_1 + \alpha_2}}(\theta) \cdot l^{\frac{\alpha_2}{\alpha_1 + \alpha_2}}(\theta). \quad (20)$$

By letting  $\alpha := \alpha_1/(\alpha_1 + \alpha_2)$ , we obtain the  $\alpha$ -pooled-posterior

$$p_g(\theta|\mathbf{y}) \propto p^\alpha(\theta) \cdot l^{1-\alpha}(\theta), \quad (21)$$

where  $0 \leq \alpha \leq 1$ .  $\square$

An engineering application of (21) is reported in social learning [17, Eq. (7)], [18, Eq. (10)]. The  $\alpha$ -pooled-posterior in (21) is a  $\alpha$ -log-linear pooling [19], [20] of the prior distribution  $p(\theta)$  and the likelihood distribution  $l(\theta)$ , which is an alternative of the linear pooling rule  $\alpha p(\theta) + (1 - \alpha)l(\theta)$ . The log-linear pooling rule is standard in Bayesian statistics for fusing multiple priors (from multiple experts) to obtain an integrated prior [19]. However, the  $\alpha$ -pooled-posterior in (21) claims that the likelihood distribution can also be used as a "prior" where the collected data  $\mathbf{y}$  serve as an expert. (But this data-driven expert does not insist on his/her opinion due to randomness of  $\mathbf{y}$ ; he/she is a flexible expert; instead, conventional non-data-driven experts are stubborn.) In addition, from (5), we can see that when  $\alpha_3 = 0$ , the entropy regularizer is removed. Therefore, it is natural to imagine that the  $\alpha$ -pooled posterior in (21) has larger entropy than the conventional Bayes' posterior  $p(\theta|\mathbf{y})$ .

From Examples 1-7, we can see that the  $(\alpha, \beta)$ -posterior, compared to the conventional Bayes' posterior (when  $\alpha = \beta = 1$ ), defines general fusing rules of the prior distribution  $p(\theta)$  and the likelihood distribution  $l(\theta)$ ; we refer to the law of defining the  $(\alpha, \beta)$ -posterior as the *generalized Bayes' rule*.

### B. Properties of Generalized Bayes' Rule

The definition of the  $(\alpha, \beta)$ -posterior motivates us to study the properties of  $\alpha$ -scaled distributions.

*Definition 3 ( $\alpha$ -Scaled Distribution):* The  $\alpha$ -scaled distribution  $h^{(\alpha)}(\theta)$  induced by the distribution  $h(\theta)$  is defined as

$$h^{(\alpha)}(\theta) = \frac{h^\alpha(\theta)}{\int h^\alpha(\theta) d\theta}, \quad (22)$$

for  $0 \leq \alpha < \infty$ , if  $\int h^\alpha(\theta) d\theta$  exists.  $\square$

As the definition implies,  $h^{(\alpha)}(\theta) \equiv h(\theta)$  for every  $\alpha$  if  $h(\theta)$  is a uniform distribution on a finite support set  $\Theta$ .

Inspired by (5), we investigate the relation between the entropy of  $h^{(\alpha)}(\theta)$  and that of  $h(\theta)$ . We start with examining the entropy of  $h^{(\alpha)}(\theta)$ .

*Theorem 2:* Suppose that  $h(\theta)$  is not a uniform distribution. The function  $\alpha \mapsto \operatorname{Ent} h^{(\alpha)}(\theta)$  is monotonically decreasing in  $\alpha$  on  $[0, \infty)$ .  $\square$

*Proof:* See Appendix C.  $\square$

Let  $E(\alpha)$  denote the entropy difference evaluated at  $\alpha \in [0, \infty)$ :

$$E(\alpha) := \operatorname{Ent} h^{(\alpha)}(\theta) - \operatorname{Ent} h(\theta). \quad (23)$$

Theorem 2 implies that  $E(\alpha)$  is monotonically decreasing in  $\alpha$  on  $[0, \infty)$ . In addition, it is obvious to see that  $E(0) > 0$  and  $E(1) = 0$ . As a result, the following corollary is immediate.

*Corollary 1:* If  $0 \leq \alpha < 1$ , then  $h^{(\alpha)}(\theta)$  has larger entropy than  $h(\theta)$ ; if  $1 < \alpha < \infty$ , then  $h^{(\alpha)}(\theta)$  has smaller entropy than  $h(\theta)$ .  $\square$

A visual illustration of the entropy difference  $E(\alpha)$  when  $h(\theta)$  is discrete is given in Fig. 1.




Fig. 1. The entropy difference  $E(\alpha) := \operatorname{Ent} h^{(\alpha)}(\theta) - \operatorname{Ent} h(\theta)$  against  $\alpha$ ;  $h(\theta)$  is a randomly generated 50-atom discrete distribution.

When  $h(\theta)$  is continuous, a concrete example for Corollary 1 is as follows.

*Example 8:* Consider an one-dimensional zero-mean Gaussian density function  $h(\theta) \propto e^{-\frac{1}{2} \frac{\theta^2}{\sigma^2}}$  where  $\sigma^2$  is the variance. The  $\alpha$ -scaled distribution is  $h^{(\alpha)}(\theta) \propto e^{-\frac{1}{2} \frac{\theta^2}{\sigma^2/\alpha}}$ . Therefore,  $h^{(\alpha)}(\theta)$  is a Gaussian density function with variance  $\sigma^2/\alpha$ . When  $0 < \alpha < 1$ , we have  $\sigma^2/\alpha > \sigma^2$  and therefore  $\operatorname{Ent} h^{(\alpha)}(\theta) > \operatorname{Ent} h(\theta)$ ; when  $\alpha > 1$ , we have  $\sigma^2/\alpha < \sigma^2$  and therefore  $\operatorname{Ent} h^{(\alpha)}(\theta) < \operatorname{Ent} h(\theta)$ . Note that  $\operatorname{Ent} h^{(\alpha)}(\theta) = \frac{1}{2} \ln(2\pi\sigma^2/\alpha) + \frac{1}{2}$ , while  $\operatorname{Ent} h(\theta) = \frac{1}{2} \ln(2\pi\sigma^2) + \frac{1}{2}$ .  $\square$

Theorem 2 and Corollary 1 collectively imply the following useful insight.

*Insight 1 (Uncertainty Awareness in Posterior):* The  $\alpha$ -scaling operation controls the entropy (i.e., the uncertainty or our trust level) of  $h^{(\alpha)}(\theta)$ . Therefore, the  $(\alpha, \beta)$ -posterior  $p_g(\theta|\mathbf{y}) \propto p^\beta(\theta) \cdot l^\alpha(\theta)$  balances the relative importance

between the prior distribution  $p(\theta)$  and the likelihood distribution  $l(\theta)$ . If we think that the likelihood (resp. prior) distribution is overly opportunistic, we use an  $\alpha$  (resp.  $\beta$ ) such that  $\text{Ent } l^{(\alpha)}(\theta)$  [resp.  $\text{Ent } p^{(\beta)}(\theta)$ ] is **increased** compared to  $\text{Ent } l(\theta)$  [resp.  $\text{Ent } p(\theta)$ ]; in contrast, if we think that the likelihood (resp. prior) distribution is overly conservative, we use an  $\alpha$  (resp.  $\beta$ ) such that  $\text{Ent } l^{(\alpha)}(\theta)$  [resp.  $\text{Ent } p^{(\beta)}(\theta)$ ] is **reduced** compared to  $\text{Ent } l(\theta)$  [resp.  $\text{Ent } p(\theta)$ ]. To be specific, if we do not trust much about the likelihood distribution (resp. the prior distribution), we use  $0 \leq \alpha < 1$  (resp.  $0 \leq \beta < 1$ ) if it is overly opportunistic, and use  $\alpha > 1$  (resp.  $\beta > 1$ ) if it is overly conservative.  $\square$

Next, we study the closeness to  $h(\theta)$  from  $h^{(\alpha)}(\theta)$ .

*Theorem 3:* The closeness to  $h(\theta)$  from  $h^{(\alpha)}(\theta)$ , i.e.,

$$\text{KL} [h(\theta) \| h^{(\alpha)}(\theta)],$$

is a monotonically increasing function if  $1 < \alpha < \infty$  and a monotonically decreasing function if  $0 \leq \alpha < 1$ . In addition, it is a convex function in  $\alpha$  on  $[0, \infty)$ .

*Proof:* See Appendix D.  $\square$

A visual illustration of Theorem 3 when  $h(\theta)$  is discrete is given in Fig. 2.




Fig. 2. A visual illustration of the closeness  $\text{KL} [h(\theta) \| h^{(\alpha)}(\theta)]$  to  $h(\theta)$  from  $h^{(\alpha)}(\theta)$ ;  $h(\theta)$  is a randomly generated 50-atom discrete distribution.

A concrete example for Theorem 3 when  $h(\theta)$  is continuous is as follows.

*Example 9:* We continue studying the setup in Example 8. As a result, we have

$$\text{KL} [h(\theta) \| h^{(\alpha)}(\theta)] = -\frac{1}{2} \ln(\alpha) + \frac{\alpha}{2} - \frac{1}{2}, \quad \alpha > 0.$$

The derivative of  $\text{KL} [h(\theta) \| h^{(\alpha)}(\theta)]$  with respect to  $\alpha$  is  $\frac{1}{2}(1 - \frac{1}{\alpha})$ , which is positive if  $\alpha > 1$  and negative if  $0 < \alpha < 1$ ; the second-order derivative is  $\frac{1}{2\alpha^2} > 0$ . Therefore,  $\text{KL} [h(\theta) \| h^{(\alpha)}(\theta)]$  is monotonically decreasing when  $0 < \alpha < 1$  and monotonically increasing when  $\alpha > 1$ . In addition,  $\text{KL} [h(\theta) \| h^{(\alpha)}(\theta)]$  is convex.  $\square$

Theorem 3 implies the following useful insight.

*Insight 2 (Level of Uncertainty):* The more  $\alpha$  deviates from 1, the farther to  $h(\theta)$  from  $h^{(\alpha)}(\theta)$ . Therefore, in  $(\alpha, \beta)$ -posterior, the more we trust the prior distribution  $p(\theta)$  [resp. the likelihood distribution  $l(\theta)$ ], the closer the value of  $\beta$  (resp.  $\alpha$ ) should be to 1.  $\square$

Insights 1 and 2 collectively suggest the usage of the  $(\alpha, \beta)$ -posterior in balancing the relative importance between the

prior knowledge and the data evidence. Next, we show another property of the  $(\alpha, \beta)$ -posterior that enables its practical usefulness. Let  $h_0(\theta)$  be the true distribution and  $h(\theta)$  the nominal distribution serving as an estimate of  $h_0(\theta)$ . The theorem below states that there exists some  $\alpha > 0$  such that the  $\alpha$ -scaled nominal distribution  $h^{(\alpha)}(\theta)$  can be closer to the true distribution  $h_0(\theta)$  than the original nominal distribution  $h(\theta)$ .

*Theorem 4:* Let  $h_0(\theta)$  be a distribution defined on  $\Theta$ . Suppose that  $h(\theta)$  is not a uniform distribution. For every  $h_0$  and  $h$  where  $h_0 \neq h$ , there exists  $\alpha \geq 0$  such that

$$\text{KL} [h_0(\theta) \| h^{(\alpha)}(\theta)] < \text{KL} [h_0(\theta) \| h(\theta)]. \quad (24)$$

*Proof:* See Appendix E.  $\square$

Theorem 4 indicates that, if a proper  $\alpha$  is given,  $h^{(\alpha)}$  can be a better (i.e., a more accurate) surrogate of  $h_0$  than  $h$ . However, the parameter  $\alpha$  cannot be theoretically specified because it depends on the underlying true but unknown distribution  $h_0(\theta)$ . The example below provides an intuitive demonstration of Theorem 4.

*Example 10:* Consider  $h_0 = [0.2, 0.8]$ ,  $h = [0.4, 0.6]$ , and  $\alpha = 2$ . We have  $h^{(2)} = [0.3, 0.7]$ , which is closer to  $h_0$  than  $h$ . For another case, consider  $h_0 = [0.4, 0.6]$ ,  $h = [0.2, 0.8]$ , and  $\alpha = 0.6$ . We have  $h^{(0.6)} = [0.3, 0.7]$ , which is closer to  $h_0$  than  $h$ . Third, we consider a continuous case. Let  $h_0(\theta) = \mathcal{N}(\theta; 0, \sigma_0^2)$  and  $h(\theta) = \mathcal{N}(\theta; 0, \sigma^2)$ . For any  $\alpha > 0$ , we have  $h^{(\alpha)}(\theta) = \mathcal{N}(\theta; 0, \sigma^2/\alpha)$ . If we choose  $\alpha := \sigma^2/\sigma_0^2$ , we have  $h^{(\alpha)}(\theta) = h_0(\theta)$ ; note that  $\alpha$  can be either larger or smaller than one.  $\square$

Theorem 4 justifies the potential that the proposed  $(\alpha, \beta)$ -posterior can outperform the conventional Bayes' rule in practice; however, the optimal parameters  $(\alpha, \beta)$  cannot be theoretically specified because they depend on true prior distribution and data-generating distribution, which are unknown.

### C. Generalized Bayes' Rule for Multiple Samples

When we have more than one sample, e.g.,  $n$  i.i.d. samples  $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$ , we can straightforwardly generalize (5) to

$$\min_{q(\theta)} \alpha_3 \text{Ent } q(\theta) + \alpha_1 \text{KL} [q(\theta) \| p(\theta)] + \frac{\alpha_2}{n} \sum_{i=1}^n \text{KL} [q(\theta) \| l_{\mathbf{y}_i}(\theta)]. \quad (25)$$

The solution of (25) is given in the corollary below.

*Corollary 2 (Multi-sample Generalized Bayes' Rule):* The generalized Bayes' posterior solving (25) for  $n$  i.i.d. samples is given by

$$p_g(\theta | \mathbf{y}_1, \dots, \mathbf{y}_n) \propto p^{\frac{\alpha_1}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta) \cdot \left[ \prod_{i=1}^n l_{\mathbf{y}_i}^{\frac{1}{n}}(\theta) \right]^{\frac{\alpha_2}{\alpha_1 + \alpha_2 - \alpha_3}}, \quad (26)$$

if  $\alpha_1 + \alpha_2 > \alpha_3$ , provided that the right-hand-side term is integrable on  $\Theta$ . When  $\alpha_1 + \alpha_2 = \alpha_3$ ,  $p_g(\theta | \mathbf{y}_1, \dots, \mathbf{y}_n)$  is an arbitrary distribution supported on the set  $\Theta^*$  where

$$\Theta^* := \arg\max_{\theta} \alpha_1 \ln p(\theta) + \alpha_2 \cdot \frac{1}{n} \sum_{i=1}^n \ln l_{\mathbf{y}_i}(\theta) \quad (27)$$

contains all weighted maximum posterior estimates; if further  $\alpha_1 = 0$ , (27) reduces to maximum likelihood estimation.  $\square$

The  $(\alpha, \beta)$ -posterior in Definition 2 under  $n$  i.i.d. samples, induced by (26), can be rewritten as

$$p_g(\theta | \mathbf{y}_1, \dots, \mathbf{y}_n) \propto p^\beta(\theta) \cdot \prod_{i=1}^n l_{\mathbf{y}_i}^{\frac{\alpha}{n}}(\theta), \quad 0 \leq \alpha, \beta \leq \infty. \quad (28)$$

Examples 1-7 can be restated accordingly; we do not repeat here.

As the data size increases (i.e.,  $n \uparrow \infty$ ), data-evidence becomes dominating and we can therefore let  $\alpha_1$  depend on  $n$  and  $\alpha_{1,n} \rightarrow 0$  as  $n \rightarrow \infty$ . An example of  $\alpha_{1,n}$  can be  $\alpha_{1,n} := 1/n$ , or  $1/n^2$  for a faster vanishing rate. As a result,  $p_g(\theta | \mathbf{y}_1, \dots, \mathbf{y}_n)$  is asymptotically equal to an  $\alpha$ -scaled likelihood distribution; cf. Example 5.

#### D. Generalized Bayes' Rule for Multiple Priors and Samples

We can further extend (25) when multiple priors are present:

$$\min_{q(\theta)} \alpha_3 \text{Ent } q(\theta) + \alpha_1 \left[ \sum_{i=1}^m \beta_i \cdot \text{KL}[q(\theta) \| p_i(\theta)] \right] + \alpha_2 \left[ \frac{1}{n} \sum_{i=1}^n \text{KL}[q(\theta) \| l_{\mathbf{y}_i}(\theta)] \right], \quad (29)$$

where  $m$  priors  $p_1(\theta), p_2(\theta), \dots, p_m(\theta)$  are available with weights  $\beta_1, \beta_2, \dots, \beta_m$ , respectively;  $\beta_i \in [0, 1]$  for every  $i \in [m]$  and  $\sum_{i=1}^m \beta_i = 1$ .

The solution of (29) is given in the corollary below.

**Corollary 3 (Multi-prior-multi-sample Generalized Bayes' Rule):** The generalized Bayes' posterior solving (29) for  $m$  priors and  $n$  i.i.d. samples is given by

$$p_g(\theta | \mathbf{y}_1, \dots, \mathbf{y}_n) \propto \left[ \prod_{i=1}^m p_i^{\beta_i}(\theta) \right]^{\frac{\alpha_1}{\alpha_1 + \alpha_2 - \alpha_3}} \cdot \left[ \prod_{i=1}^n l_{\mathbf{y}_i}^{\frac{1}{n}}(\theta) \right]^{\frac{\alpha_2}{\alpha_1 + \alpha_2 - \alpha_3}}, \quad (30)$$

if  $\alpha_1 + \alpha_2 > \alpha_3$ , provided that the right-hand-side term is integrable on  $\Theta$ . When  $\alpha_1 + \alpha_2 = \alpha_3$ ,  $p_g(\theta | \mathbf{y}_1, \dots, \mathbf{y}_n)$  is an arbitrary distribution supported on the set  $\Theta^*$  where

$$\Theta^* := \underset{\theta}{\operatorname{argmax}} \alpha_1 \sum_{i=1}^m \beta_i \cdot \ln p_i(\theta) + \alpha_2 \sum_{i=1}^n \frac{1}{n} \cdot \ln l_{\mathbf{y}_i}(\theta) \quad (31)$$

contains all weighted maximum posterior estimates.  $\square$

The  $(\alpha, \beta)$ -posterior in Definition 2 under  $m$  priors and  $n$  i.i.d. samples, induced by (30), can be rewritten as

$$p_g(\theta | \mathbf{y}_1, \dots, \mathbf{y}_n) \propto \prod_{i=1}^m p_i^{\beta \cdot \beta_i}(\theta) \cdot \prod_{i=1}^n l_{\mathbf{y}_i}^{\frac{\alpha}{n}}(\theta), \quad 0 \leq \alpha, \beta \leq \infty. \quad (32)$$

Examples 1-7 can be restated accordingly; we do not repeat here.

#### E. Illustrating Examples

In this subsection, we provide some illustrating examples of the  $\alpha$ -scaled distribution and the  $(\alpha, \beta)$ -posterior. First, we visualize the  $\alpha$ -scaled distribution  $h^{(\alpha)}(\theta)$  when  $h(\theta)$  is discrete; see Fig. 3. As we can see, when  $0 \leq \alpha < 1$ , the probabilities become more balanced, while when  $\alpha > 1$ , they become more unbalanced.




Fig. 3. A 3-atom discrete distribution  $h(y)$  and induced  $h^{(\alpha)}(y)$  with different  $\alpha \in \{0.5, 2\}$ . Under  $\alpha = 0.5$ ,  $\text{Ent } h^{(\alpha)}(y) > \text{Ent } h(y)$  (i.e., the former has more balanced masses, while the latter has more unbalanced masses). Under  $\alpha = 2$ ,  $\text{Ent } h^{(\alpha)}(y) < \text{Ent } h(y)$ .

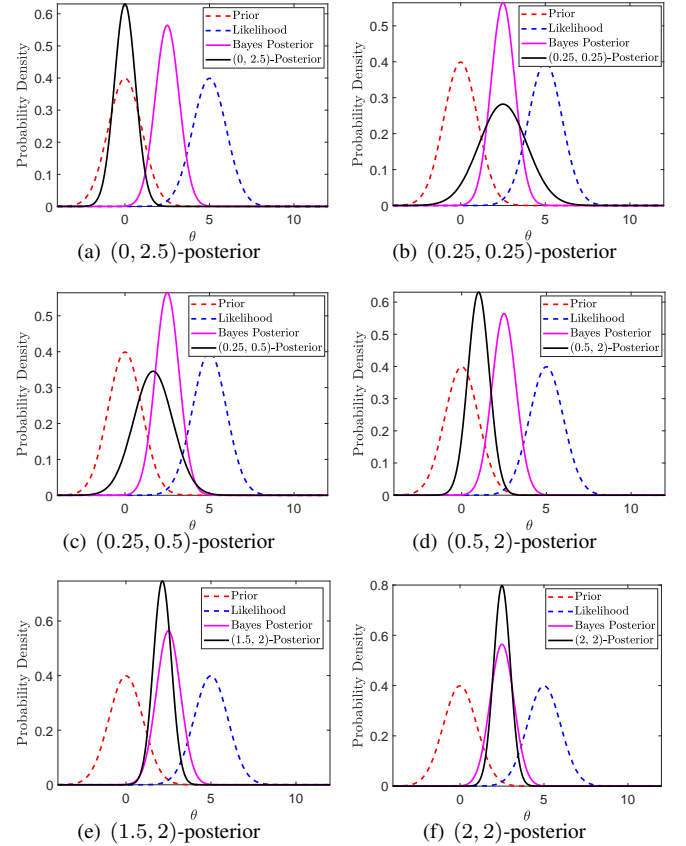


Fig. 4. Illustrating examples of  $(\alpha, \beta)$ -posterior. The prior distribution is  $p(\theta) := \mathcal{N}(\theta; 0, 1)$ . The likelihood distribution is  $l_{\mathbf{y}:=5}(\theta) := \mathcal{N}(\theta; 5, 1)$ .

Next, we visualize the  $(\alpha, \beta)$ -posterior, compared with the conventional Bayes' posterior. We focus on a mean estimation problem where  $\theta_0$  is the true mean of the random variable  $y$ . We suppose that the prior distribution is  $p(\theta) := \mathcal{N}(\theta; 0, 1)$ .



The likelihood function is  $\theta \mapsto \mathcal{N}(y; \theta, 1)$ ; if we assume that the measurement is  $y := 5$ , the likelihood distribution is therefore  $l(\theta) := \mathcal{N}(\theta; 5, 1)$ . The visualizations of the  $(\alpha, \beta)$ -posterior, under different value pairs of  $(\alpha, \beta)$ , are given in Fig. 4.

As motivational examples, we only interpret the first three sub-figures. In Fig. 4(a),  $\alpha = 0$  means that the likelihood distribution is absolutely unreliable and the data evidence is completely ignored. Therefore, the  $(0, 2.5)$ -posterior distribution tends to overlap with the prior distribution. Moreover, because  $\beta = 2.5 > 1$ , the specified prior distribution is thought to be overly conservative and we propose to increase its concentration. In Fig. 4(b),  $\alpha = \beta = 0.5$  means that both the prior distribution and the likelihood distribution are believed to be overly opportunistic, and therefore, we propose to reduce their concentration. As a result, the  $(0.25, 0.25)$ -posterior distribution has larger entropy than the conventional Bayes' posterior distribution. In Fig. 4(c), since  $\alpha = 0.25 < 0.5 = \beta$ , both the prior distribution and the likelihood distribution are thought to be opportunistic, but the prior distribution is thought to be less opportunistic than the likelihood distribution.

#### F. Examples of Application

The generalized (or uncertainty-aware) Bayes' rule (9), i.e., the  $(\alpha, \beta)$ -posterior, has several potential applications in statistical machine learning and statistical signal processing. We specifically discuss classification and estimation problems. Suppose that the true joint distribution is  $p_0(\mathbf{y}, \boldsymbol{\theta})$ , the true prior distribution is  $p_0(\boldsymbol{\theta})$ , and the true likelihood distribution is  $l_{0,\mathbf{y}}(\boldsymbol{\theta})$  under the measurement  $\mathbf{y}$ . Let the nominal prior distribution and the nominal likelihood distribution be  $p(\boldsymbol{\theta})$  and  $l_{\mathbf{y}}(\boldsymbol{\theta})$ , respectively.

1) *Bayesian Classification*: Classification is a fundamental component of statistical machine learning. Let  $\Theta := \{1, 2, 3, \dots, c\}$  be a categorical set where the integer  $c$  is finite. The Bayes classifier based on the conventional Bayes' rule is

$$\hat{\theta}(\mathbf{y}) := \operatorname{argmax}_{\theta \in \Theta} \log p_0(\theta) + \log l_{0,\mathbf{y}}(\theta). \quad (33)$$

In practice, however, the true distributions  $p_0(\theta)$  and  $l_{0,\mathbf{y}}(\theta)$  are unknown. By employing the nominal distributions  $p(\theta)$  and  $l_{\mathbf{y}}(\theta)$ , the generalized Bayes classifier based on the  $(\alpha, \beta)$ -posterior (9) can be obtained as

$$\hat{\theta}_g(\mathbf{y}) := \operatorname{argmax}_{\theta \in \Theta} \beta \log p(\theta) + \alpha \log l_{\mathbf{y}}(\theta), \quad (34)$$

which is equivalent, in the sense of the same classifier, to

$$\hat{\theta}_g(\mathbf{y}) := \operatorname{argmax}_{\theta \in \Theta} (1 - \lambda) \log p(\theta) + \lambda \log l_{\mathbf{y}}(\theta), \quad (35)$$

for  $\lambda := \alpha/(\alpha + \beta)$ . As we can see, for classification problems, the  $(\alpha, \beta)$ -posterior with  $\alpha, \beta \in [0, \infty)$  has the same effect as the  $(\lambda, 1 - \lambda)$ -posterior with  $\lambda \in [0, 1]$ ; cf. Example 7; only the ratio  $\alpha/\beta$  matters.

2) *Bayesian Estimation*: Estimation aims to infer an unknown parameter  $\boldsymbol{\theta}$  based on measured data  $\mathbf{y}$ . It is of particular importance in statistical signal processing. Upon

collecting  $\mathbf{y}$ , the Bayes estimator based on the conventional Bayes's rule is

$$\hat{\theta}(\mathbf{y}) := \int \boldsymbol{\theta} p_0(\boldsymbol{\theta}|\mathbf{y}) d\boldsymbol{\theta}, \quad (36)$$

where  $p_0(\boldsymbol{\theta}|\mathbf{y}) \propto p_0(\boldsymbol{\theta}, \mathbf{y})$  is the true Bayes' posterior distribution. In practice, however, the true distributions  $p_0(\boldsymbol{\theta})$  and  $l_{0,\mathbf{y}}(\boldsymbol{\theta})$  are unknown. By employing the nominal distributions  $p(\boldsymbol{\theta})$  and  $l_{\mathbf{y}}(\boldsymbol{\theta})$ , the generalized Bayes estimator based on the  $(\alpha, \beta)$ -posterior (9) can be obtained as

$$\hat{\theta}_g(\mathbf{y}) := \int \boldsymbol{\theta} p_g(\boldsymbol{\theta}|\mathbf{y}) d\boldsymbol{\theta}. \quad (37)$$

3) *Other Examples*: In this subsection, we discuss other applications of the  $(\alpha, \beta)$ -posterior. First, we give an example of application in statistical signal processing.

*Example 11 (Uncertainty-Aware Particle Filter)*: Suppose that  $\Theta$  is a finite discrete set containing  $r$  points (each point is known as a particle): i.e.,  $\Theta := \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_r\}$ . For each particle  $\boldsymbol{\theta}_i, i \in [r]$ , the likelihood evaluated at the measurement  $\mathbf{y}$  is assumed to be  $p(\mathbf{y}|\boldsymbol{\theta}_i)$ . Further, we suppose that the prior distribution is  $\mathbf{p} := [p_1, p_2, \dots, p_r]$  and the  $\mathbf{y}$ -likelihood distribution (recall Definition 1) is  $\mathbf{l} := [l_1, l_2, \dots, l_r]$ . Then the generalized Bayes' posterior, i.e., the  $(\alpha, \beta)$ -posterior, of particles is equal to

$$p_g(\boldsymbol{\theta}_i|\mathbf{y}) \propto p_i^\beta \cdot l_i^\alpha, \quad \forall i \in [r].$$

The above formula leads to the uncertainty-aware particle filter for dynamic stochastic nonlinear systems.  $\square$

Second, we give an example of application in statistical machine learning. It is shown that the uncertainty-aware Maximum *A-Posteriori* (MAP) estimation can lead to the popular ridge regression.

*Example 12 (Uncertainty-Aware MAP Estimation)*: We consider a nonlinear regression model  $y = f(\mathbf{x}; \boldsymbol{\theta}) + \epsilon$  where  $y \in \mathbb{R}$  is the response,  $\mathbf{x} \in \mathbb{R}^m$  is the feature vector,  $\epsilon \sim \mathcal{N}(0, 1)$  is the regression error, and  $\boldsymbol{\theta}$  is the parameter vector. Supposing the prior distribution of  $\boldsymbol{\theta}$  is  $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ , upon the collection of the data  $(y, \mathbf{x})$ , the MAP estimator of  $\boldsymbol{\theta}$  is

$$\min_{\boldsymbol{\theta} \in \Theta} [y - f(\mathbf{x}; \boldsymbol{\theta})]^2 + \boldsymbol{\theta}^\top \boldsymbol{\theta}.$$

By using the  $(\alpha, \beta)$ -posterior rule, the uncertainty-aware MAP estimator of  $\boldsymbol{\theta}$  can be written as

$$\min_{\boldsymbol{\theta} \in \Theta} \alpha [y - f(\mathbf{x}; \boldsymbol{\theta})]^2 + \beta \boldsymbol{\theta}^\top \boldsymbol{\theta}.$$

By letting  $\lambda := \beta/\alpha$ , we have the  $\lambda$ -ridge regression

$$\min_{\boldsymbol{\theta} \in \Theta} [y - f(\mathbf{x}; \boldsymbol{\theta})]^2 + \lambda \boldsymbol{\theta}^\top \boldsymbol{\theta}.$$

Therefore, the popular ridge regression method in statistical machine learning can be interpreted as an uncertainty-aware MAP estimation.  $\square$

Third, we discuss an example of application that is popular in both statistical signal processing and statistical machine learning.

*Example 13 (Uncertainty-Aware Bayesian Model Averaging)*: Suppose that we have  $r$  models to account for a signal

processing or a machine learning problem. Let the prior distribution of models be  $\mathbf{p} := [p_1, p_2, \dots, p_r]$  and the likelihood distribution (recall Definition 1) be  $\mathbf{l} := [l_1, l_2, \dots, l_r]$ . Then the generalized Bayes' posterior, i.e., the  $(\alpha, \beta)$ -posterior, of models is equal to

$$p_g(\theta_i | \mathbf{y}) \propto p_i^\beta \cdot l_i^\alpha, \quad \forall i \in [r].$$

The above formula induces the uncertainty-aware Bayesian model averaging method.  $\square$

### G. Hyper-Parameter Tuning

This subsection discusses the determination methods for the hyper-parameters  $(\alpha, \beta)$  in practice. The main purpose is to find some  $(\alpha, \beta)$  such that the  $(\alpha, \beta)$ -posterior can outperform the conventional Bayes' posterior; cf. Theorem 4. Suppose that the collected history data set is  $\mathcal{D}_t := \{(\theta_1, \mathbf{y}_1), (\theta_2, \mathbf{y}_2), \dots, (\theta_t, \mathbf{y}_t)\}$ .

Let  $\omega := [\alpha, \beta]^\top$  denote the parameter vector of an  $(\alpha, \beta)$ -posterior and  $f(\omega)$  the loss function of using the  $(\alpha, \beta)$ -posterior for a specific application, where  $f : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ . We assume that  $\omega$  takes values on  $[0, \tau]^2$  where  $\tau > 1$  is a bounded positive real number (for example,  $\tau = 5$ ). The optimal parameter tuning for  $\omega$  can be formulated as a minimization problem

$$\min_{\omega \in [0, \tau]^2} f(\omega). \quad (38)$$

Examples of  $f(\omega)$  are given below.

*Example 14 (Classification):* For classification problems, the loss function is the misclassification probability, i.e.,

$$f(\omega) := \int \int \mathbb{I}_{\{\theta \neq \hat{\theta}_{g, \omega}(\mathbf{y})\}} \cdot p_0(\mathbf{y}, \theta) d\mathbf{y} d\theta, \quad (39)$$

where  $\mathbb{I}_{\{\cdot\}}$  denotes the indicator function and the generalized classifier  $\hat{\theta}_{g, \omega}(\mathbf{y})$  is defined in (34). Note that  $\hat{\theta}_g(\mathbf{y})$  in (34) depends on  $\omega$ . If the classification method in (35) is employed, we have the cost function

$$f(\lambda) := \int \int \mathbb{I}_{\{\theta \neq \hat{\theta}_{g, \lambda}(\mathbf{y})\}} \cdot p_0(\mathbf{y}, \theta) d\mathbf{y} d\theta, \quad \lambda \in [0, 1]$$

since the classifier defined in (35) is parameterized by  $\lambda$ .  $\square$

*Example 15 (Estimation):* For estimation problems, the loss function in the mean squared error (MSE) sense can be written as

$$f(\omega) := \int \int [\theta - \hat{\theta}_{g, \omega}(\mathbf{y})]^\top [\theta - \hat{\theta}_{g, \omega}(\mathbf{y})] \cdot p_0(\mathbf{y}, \theta) d\mathbf{y} d\theta, \quad (40)$$

where the generalized estimator  $\hat{\theta}_{g, \omega}(\mathbf{y})$  is defined in (37). Note that  $\hat{\theta}_g(\mathbf{y})$  in (37) depends on  $\omega$ .  $\square$

In practice, however, the loss function  $f(\omega)$  is unknown due to the unavailability of  $p_0(\mathbf{y}, \theta)$ . Based on the historical data  $\mathcal{D}_t$ , we can use data-driven estimate  $\hat{f}(\omega)$  as a surrogate of  $f(\omega)$ . For instance,

$$\hat{f}(\omega) = \frac{1}{t} \sum_{i=1}^t \mathbb{I}_{\{\theta_i \neq \hat{\theta}_{g, \omega}(\mathbf{y}_i)\}} \quad (41)$$

in Example 14 and

$$\hat{f}(\omega) = \frac{1}{t} \sum_{i=1}^t [\theta_i - \hat{\theta}_{g, \omega}(\mathbf{y}_i)]^\top [\theta_i - \hat{\theta}_{g, \omega}(\mathbf{y}_i)] \quad (42)$$

in Example 15. When the training data set  $\mathcal{D}_t$  is sufficiently large,  $\hat{f}(\omega)$  can be a point-wisely good estimate of  $f(\omega)$ .

In the following, we suggest two methods for tuning  $\omega$ . Since (38) is a low-dimensional optimization problem with at most two variables on a hyper-cube domain [cf. (35) and (37)], both methods can be statistically and computationally efficient; for empirical validation, see experiments in Section III.

*1) Grid Search:* For simplicity in operation, we can generate a two-dimensional uniform grid  $\Omega_{\text{grid}}$  on  $[0, \tau]^2$  to obtain discrete empirical evaluations  $\hat{f}(\omega)$  for  $\omega \in \Omega_{\text{grid}} \subset [0, \tau]^2$ ; note that  $(1, 1)$  should be included in  $\Omega_{\text{grid}}$ . Then, we solve

$$(\alpha^*, \beta^*) = \underset{\omega \in \Omega_{\text{grid}}}{\operatorname{argmin}} \hat{f}(\omega) \quad (43)$$

to obtain the best parameter pair  $(\alpha^*, \beta^*)$ . Since  $(1, 1) \in \Omega_{\text{grid}}$ ,  $\hat{f}([\alpha^*, \beta^*]^\top) \leq \hat{f}([1, 1]^\top)$ . Namely, the loss under the proposed  $(\alpha, \beta)$ -posterior is no larger than that under the usual Bayes' posterior. This can empirically but sufficiently validate the advantage of the  $(\alpha, \beta)$ -posterior over the conventional Bayes' posterior.

*2) Surrogate Optimization:* The grid search method cannot pursue the global optimality on  $[0, \tau]^2$ . Considering the optimization problem (38) and its characteristics (i.e.,  $f$  is unknown but some noisy evaluations  $\hat{f}$  are available), we can leverage the surrogate optimization frameworks, for example, radial-basis-function surrogate optimization [21], Gaussian-process surrogate optimization (i.e., Bayesian optimization) [22], [23]. To put it simply, surrogate optimization uses as few as possible noisy evaluations  $\hat{f}(\omega)$  to learn the true function  $f(\omega)$  in an online manner, and simultaneously search for the globally optimal minimizer(s)  $\omega^*$  such that  $f(\omega^*) \leq f(\omega)$  for all  $\omega$  on  $[0, \tau]^2$ .

## III. CONCRETE APPLICATIONS

The generalized (or uncertainty-aware) Bayes' rule (9) can give birth to several uncertainty-aware Bayesian machine learning methods and uncertainty-aware Bayesian signal processing methods. As demonstrations and without loss of generality, in this section, we focus on naive Bayes classification problems in machine learning and state estimation (i.e., state filtering) problems of dynamic stochastic state-space systems in signal processing. The main purpose is to show the existence of  $(\alpha, \beta)$  such that the  $(\alpha, \beta)$ -posterior can outperform the conventional Bayes' posterior; cf. Theorem 4. All the source data and codes are available online at GitHub: <https://github.com/Sprtm-Asleaf/Bayes-Rule>. The experimental results are obtained by a Lenovo laptop with 16G RAM and 11th Gen Intel(R) Core(TM) i5-11300H CPU @ 3.10GHz.

### A. Bayesian Classification: Text and Image Classification

We consider two real-world classification problems. One is a natural-language text classification problem, while the other is a medical image classification problem.



1) *IMDB Dataset*: We investigate the performances of the generalized Bayes classifier in (35) on the IMDB dataset containing movie reviews [24, Sec. 4.3.2]. This problem is also known as text-based sentiment analysis where we determine whether a review of a movie is positive or negative. The multinomial naive Bayes classification algorithm with Laplacian smoothing is employed; that is, the features are distributed according to multinomial distributions because movie reviews are typically represented as word vector counts, and the features are assumed to be conditionally independent given the target class; see [25].

For every sample size  $L$  in  $\{50, 100, 250, 500, 1000, 2000, 5000, 10000\}$ , we conduct 500 independent Monte–Carlo tests. In detail, in each Monte–Carlo trial, we randomly select  $L$  samples from the IMDB dataset, of which 80% are training samples and 20% are testing samples. For each Monte–Carlo test, the radial-basis-function surrogate optimization on  $[0, 1]$  is employed to find the empirically optimal hyper-parameter  $\lambda$  in (35);<sup>5</sup> the testing performance, which is a function of  $\lambda$ , serves as the loss function in (38) to be minimized. The surrogate-optimization-based search process starts with  $\lambda = 0.5$  (i.e., the conventional Bayes classifier) so that the performance of the generalized Bayes classifier (35) is at least as good as that of the conventional Bayes classifier. The experimental results are shown in Fig. 5.

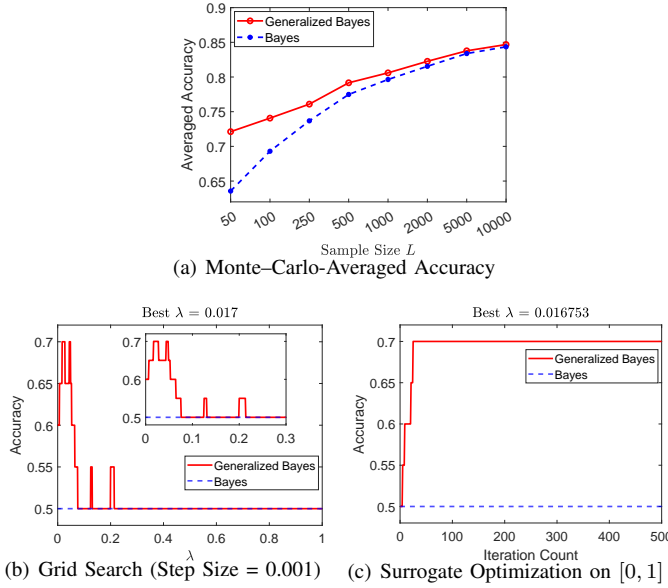


Fig. 5. Experimental results on the IMDB data set. (a): The classification accuracy against sample size  $L$ , averaged over 500 Monte–Carlo episodes. (b): The grid search for  $\lambda$  with step size of 0.001 in a Monte–Carlo trial when  $L = 100$ . (c): The visual illustration of the surrogate optimization process for  $\lambda$  on  $[0, 1]$ ; the found optimal value is  $\lambda = 0.016753$ .

As sample size  $L$  becomes larger, the nominal distributions  $p(\theta)$  and  $l_{\mathbf{y}}(\theta)$  get closer to the true distributions  $p_0(\theta)$  and  $l_{0,\mathbf{y}}(\theta)$ , respectively. Therefore, the conventional Bayes classifier tends to obtain the best classification accuracy as  $L$  increases; cf. Fig. 5(a). In Figs. 5(b) and 5(c), the parameter tuning process is visualized for a Monte–Carlo trial when

<sup>5</sup>See MATLAB’s `surrogateopt` function at <https://mathworks.com/help/gads/surrogateopt.html>.

$L = 100$ . Since the found optimal  $\lambda = 0.016753$ , any hyper-parameter pair of  $(\alpha, \beta)$  in (34) such that  $\alpha/(\alpha + \beta) = 0.016753$  is optimal. Both two tuning methods find the best accuracy of 0.7. The average running times of the two tuning methods, against sample size  $L$ , are shown in Fig. 6. Note that the larger the testing sample size, the more running time is needed to calculate the empirical performance (41).




Fig. 6. Average running times of the two tuning methods against sample size  $L$ . The time difference of the two tuning methods remains about 1.5 seconds for every  $L$ . The running time of grid search depends on the step size (NB: the smaller the step size, the larger the maximum accuracy, but the slower the running speed), while that of surrogate optimization depends on the maximum allowed iteration count (NB: the larger, the larger, but the slower).

2) *UCI Breast Cancer Dataset*: We examine the performances of the generalized Bayes classifier in (35) on the UCI breast cancer dataset [26]. This is an image-classification-based medical diagnosis problem. The Gaussian naive Bayes classification algorithm is employed; that is, the features are distributed according to Gaussian distributions, and the features are assumed to be conditionally independent given the target class; see [25].

Since there are only 569 instances in this data set, we use all of them in the experiment. We conduct 500 independent Monte–Carlo tests. In each Monte–Carlo trial, 80% of the 569 instances are randomly drawn to serve as training samples, while the remaining 20% are testing samples. The radial-basis-function surrogate optimization on  $[0, 1]$  is employed to find the empirically optimal hyper-parameter  $\lambda$  in (35); the testing performance, which is a function of  $\lambda$ , serves as the loss function in (38) to be minimized. The surrogate-optimization-based search process starts with  $\lambda = 0.5$  (i.e., the conventional Bayes classifier) so that the performance of the generalized Bayes classifier (35) is at least as good as that of the conventional Bayes classifier. The experimental results are shown in Fig. 7(a). In a Monte–Carlo trial, both parameter tuning methods find the best accuracy of 0.9027; see Figs. 7(b) and 7(c).

## B. Bayesian Estimation: State Estimation

From the perspective of applied statistics, state estimation problems can be interpreted as sequential Bayesian inference. Briefly speaking, state estimation aims to estimate the unknown (i.e., unobservable) state vector  $\mathbf{x}_k$  at the discrete time step  $k$  using the known (i.e., observable) measurement set  $\mathcal{Y}_k := (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_k)$  and the probabilistic model  $p_{\mathbf{x}_k, \mathcal{Y}_k}(\mathbf{x}_k, \mathbf{Y}_k)$ ; the probabilistic model  $p_{\mathbf{x}_k, \mathcal{Y}_k}(\mathbf{x}_k, \mathbf{Y}_k)$  is induced by the stochastic state-space models. The aim of state estimation is to obtain the posterior state distribution  $p_{\mathbf{x}_k | \mathcal{Y}_k}(\mathbf{x}_k | \mathbf{Y}_k)$  or the posterior mean  $\mathbb{E}(\mathbf{x}_k | \mathcal{Y}_k)$ .

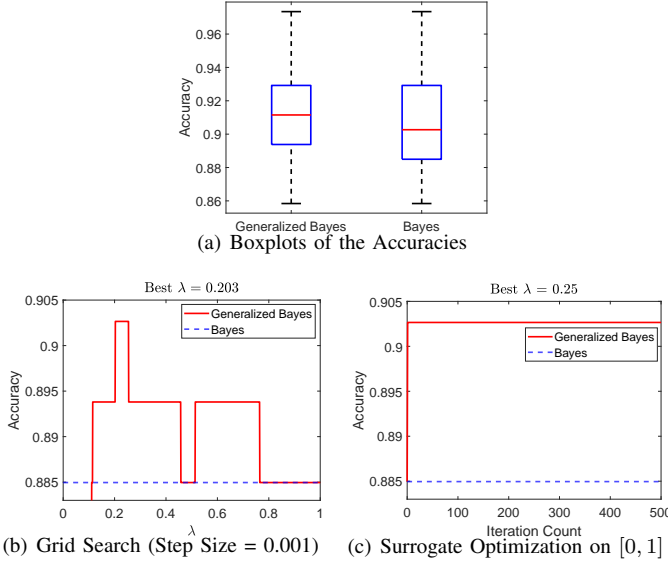


Fig. 7. Experimental results on the UCI Breast Cancer data set. (a): The boxplots for the classification accuracies of 500 Monte-Carlo episodes. (b): The grid search for  $\lambda$  with step size of 0.001 in a Monte-Carlo trial. (c): The visual illustration of the surrogate optimization process for  $\lambda$  on  $[0, 1]$ ; the found optimal value is  $\lambda = 0.25$ . (**Running Times:** The average running time for the grid search method is 0.2 seconds and for the surrogate optimization method is 1.5 seconds.)

1) *Uncertainty-Aware Kalman Filter:* We consider the state estimation problem of linear Gaussian state-space models [27, Chap. 3]

$$\begin{cases} \mathbf{x}_k = \mathbf{F}_{k-1}\mathbf{x}_{k-1} + \mathbf{G}_{k-1}\mathbf{w}_{k-1}, \\ \mathbf{y}_k = \mathbf{H}_k\mathbf{x}_k + \mathbf{v}_k, \end{cases} \quad (44)$$

where, for every time  $k$ , the system matrices  $\mathbf{F}_k$ ,  $\mathbf{G}_k$ , and  $\mathbf{H}_k$  are assumed to be known. The process noise vector and the measurement noise vector are denoted by  $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$  and  $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k)$ , respectively, and  $\mathbf{Q}_k$  and  $\mathbf{R}_k$  are assumed to be known, for every  $k$ . For this linear Gaussian system, under several uncorrelatedness assumptions among  $\mathbf{x}_0$ ,  $\{\mathbf{w}_k\}_{\forall k}$ , and  $\{\mathbf{v}_k\}_{\forall k}$  (see [27, p. 38]), the Kalman filter is the optimal state estimator in the sense of minimum mean-squared error.

Suppose that at time  $k-1$ , the posterior state distribution is

$$\hat{\mathbf{x}}_{k-1|k-1} \sim \mathcal{N}(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{P}_{k-1|k-1}).$$

At time  $k$ , the prior state distribution is

$$\hat{\mathbf{x}}_{k|k-1} \sim \mathcal{N}(\hat{\mathbf{x}}_{k|k-1}, \mathbf{P}_{k|k-1})$$

where  $\mathbf{P}_{k|k-1} := \mathbf{F}_{k-1}\mathbf{P}_{k-1|k-1}\mathbf{F}_{k-1}^\top + \mathbf{G}_{k-1}\mathbf{Q}_{k-1}\mathbf{G}_{k-1}^\top$  and  $\hat{\mathbf{x}}_{k|k-1} := \mathbf{F}_{k-1}\hat{\mathbf{x}}_{k-1|k-1}$ ; the measurement distribution conditioned on state  $\mathbf{x}_k$  is

$$\mathbf{y}_k|\mathbf{x}_k \sim \mathcal{N}(\mathbf{H}_k\mathbf{x}_k, \mathbf{R}_k).$$

Therefore, by applying the  $(\alpha, \beta)$ -posterior rule (9), the prior state distribution should be modified to

$$\hat{\mathbf{x}}_{k|k-1} \sim \mathcal{N}(\hat{\mathbf{x}}_{k|k-1}, \mathbf{P}_{k|k-1}/\beta)$$

and the conditional measurement distribution should be modified to

$$\mathbf{y}_k|\mathbf{x}_k \sim \mathcal{N}(\mathbf{H}_k\mathbf{x}_k, \mathbf{R}_k/\alpha).$$

Note that for a Gaussian distribution  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ , the  $\alpha$ -scaled version is equal to  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}/\alpha)$ . As a result, the  $(\alpha, \beta)$ -uncertainty-aware Kalman filter can be accordingly obtained by just applying the following two assignment operations in each Kalman iteration:  $\mathbf{P}_{k|k-1} \leftarrow \mathbf{P}_{k|k-1}/\beta$  and  $\mathbf{R}_k \leftarrow \mathbf{R}_k/\alpha$ , for  $\alpha, \beta > 0$ .

This modification rule is reminiscent of the distributionally robust state estimation (DRSE) method proposed in [28]; see also [10, p. 22, p. 53] for intuitive interpretations. Hence, when  $0 < \alpha, \beta \leq 1$ , the  $(\alpha, \beta)$ -uncertainty-aware Kalman filter is equivalent to distributionally robust state estimation methods in [28], [29], provided that there are no outliers in measurements. Since  $0 < \alpha, \beta \leq 1$ ,  $\mathbf{P}_{k|k-1}$  and  $\mathbf{R}_k$  are assumed to be overly opportunistic, and therefore, they need to be inflated. When  $\alpha, \beta \geq 1$ , the values of  $\mathbf{P}_{k|k-1}$  and  $\mathbf{R}_k$  are reduced, which implies that  $\mathbf{P}_{k|k-1}$  and  $\mathbf{R}_k$  are assumed to be overly conservative. However, from the technical derivations of the DRSE method, the value reduction of  $\mathbf{P}_{k|k-1}$  and  $\mathbf{R}_k$  cannot be achieved in the DRSE method. In this sense, therefore, the  $(\alpha, \beta)$ -uncertainty-aware Kalman filter presented in this paper generalizes the DRSE method in [28], [29] when there are no measurement outliers: the former can handle not only opportunistic cases (by inflating covariances) but also conservative cases (by reducing covariances), while the latter can only address opportunistic cases. (NB: To combat opportunism is to introduce robustness/conservatism; i.e., robust methods innately cannot fight against conservatism.)

The experiments of the DRSE method in [10, Chap. 2] can sufficiently support the practical values of the  $(\alpha, \beta)$ -uncertainty-aware Kalman filter; just notice the mathematical equivalence to the DRSE method (when  $0 < \alpha, \beta \leq 1$ ). We do not repeat these experiments here.

2) *Uncertainty-Aware Particle Filter:* The particle filter is standard to handle the state estimation problem of nonlinear systems [30]. By employing the  $(\alpha, \beta)$ -posterior in (9), the  $(\alpha, \beta)$ -uncertainty-aware particle filter can be straightforwardly obtained; see Example 11.

As an illustration, we specifically consider the state estimation problem of a one-dimensional nonlinear system model [16], [30], [31]

$$\begin{cases} \mathbf{x}_k = \frac{\mathbf{x}_{k-1}}{2} + \frac{25\mathbf{x}_{k-1}}{1 + \mathbf{x}_{k-1}^2} + 8\cos(1.2k) + \mathbf{w}_{k-1}, \\ \mathbf{y}_k = \frac{\mathbf{x}_k^2}{20} + 0.5\sin(\mathbf{x}_k) + \mathbf{v}_k, \end{cases} \quad (45)$$

where for every  $k$ ,  $\mathbf{w}_k \sim \mathcal{N}(0, 10)$  and  $\mathbf{v}_k \sim \mathcal{N}(0, 1)$ ; for every  $k_1$  and  $k_2$ ,  $\mathbf{w}_{k_1}$  and  $\mathbf{w}_{k_2}$  are uncorrelated, so are  $\mathbf{v}_{k_1}$  and  $\mathbf{v}_{k_2}$ ; for every  $k$ ,  $\mathbf{w}_k$  and  $\mathbf{v}_k$  are uncorrelated. As in [16], we assume that the *nominal* measurement equation is

$$\mathbf{y}_k = \frac{\mathbf{x}_k^2}{20} + \mathbf{v}_k,$$

which is slightly different from the true one in (45). Therefore, in this case, there is a modeling uncertainty in the measurement equation, i.e., in likelihood distributions at every  $k$ . As a result, to combat this model uncertainty and improve the filtering accuracy, we should use  $\beta = 1$  and  $\alpha < 1$ .

For the purpose of experimental demonstration, we use 50, 100, and 200 particles, respectively, to report the performance

of the  $(\alpha, 1)$ -uncertainty-aware particle filter. The systematic resampling method is adopted to address particle degeneracy, and the effective sample size is set to half the number of particles. Given the number of particles, the experiment is independently conducted with 500 episodes and each episode contains 100 time steps. The performance measure, in every episode, is the rooted time-averaged mean-squared error (RTAMSE) along 100 time steps, i.e.,

$$\sqrt{\frac{1}{100} \sum_{k=1}^{100} (x_k - \hat{x}_k)^2}$$

where  $\hat{x}_k$  denotes the estimate of the true value  $x_k$ ; the overall performance measure for a given number of particles is the averaged RTAMSEs of the 500 episodes. The filtering result of the  $(\alpha, 1)$ -uncertainty-aware particle filter is shown in Fig. 8, plotted against the value of  $\alpha$ .




Fig. 8. Averaged-RTAMSE performance of the  $(\alpha, 1)$ -uncertainty-aware particle filter, against the value of  $\alpha$ . The larger the number of particles, the smaller the averaged RTAMSE. For a given number of particles, the  $(\alpha, 1)$ -uncertainty-aware particle filter with  $\alpha \approx 0.25$  works best.

As we can see from Fig. 8, when  $\alpha < 1$  but  $\alpha$  is not close to zero, the  $(\alpha, 1)$ -uncertainty-aware particle filter can outperform the conventional particle filter: This is because when there exists model uncertainty in the measurement equation, the likelihood distribution of particles at each time step tends to be unreliable, and therefore, we need to reduce the concentration (i.e., reduce the trust level and improve the entropy) of the likelihood distribution to cope with the uncertainty. In addition, the results suggest that the superiority of the  $(\alpha, 1)$ -uncertainty-aware particle filter tends to be more significant as the number of particles decreases: This implies that the  $(\alpha, 1)$ -uncertainty-aware particle filter has the innate ability to fight against the particle degeneracy issue. When  $\alpha$  is overly small (i.e., close to zero), the  $(\alpha, 1)$ -uncertainty-aware particle filter almost ignores information from the measurements, and therefore, the filter diverges: This indicates that uncertain information is at least better than no information.

### 3) Uncertainty-Aware Interactive Multiple Model Filter:

The interactive multiple model (IMM) filter is standard to handle the state estimation problem of jump linear systems [32], [33]. By employing the  $(\alpha, \beta)$ -posterior in (9), the  $(\alpha, \beta)$ -uncertainty-aware IMM filter can be straightforwardly obtained. One may just imagine the models' prior weights as a prior distribution and the models' likelihoods (evaluated at a given measurement) as a likelihood distribution; the aim is to infer the models' posterior distribution and then compute the posterior state estimate; see Example 13.

As an illustration, we consider two real-world one-dimensional multi-model target tracking problems; as claimed in [34], focusing on only one coordinate does not lose the generality because the motions of a dynamic target in different coordinates can be independently tracked. Mathematically, it is a state estimation problem of a jump linear system model [35]–[37]

$$\begin{cases} \mathbf{x}_k = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \mathbf{x}_{k-1} + \begin{bmatrix} T^2/2 \\ T \end{bmatrix} [a_{j_{k-1}, k-1} + w_{k-1}] \\ y_k = p_k + v_k, \end{cases}$$

where the state vector  $\mathbf{x}_k$  is defined as  $\mathbf{x}_k := [p_k, s_k]^T$ ;  $p_k \in \mathbb{R}$  denotes the position,  $s_k \in \mathbb{R}$  the speed, and  $a_{j_k, k} \in \mathbb{R}$  the maneuvering acceleration of a moving target at time  $k$ ; the positive-integer-valued discrete random variable  $j_k$  denotes the system's operating mode at time  $k$ ;  $T > 0$  is the sampling time between two discrete time indices;  $w_k \in \mathbb{R}$  is the acceleration modeling noise and  $v_k \in \mathbb{R}$  the sensor's observation noise at time  $k$ . At time  $k$ , the maneuvering acceleration may randomly take any one of the following values

$$a_{j_k, k} = \begin{cases} 0, & j_k = 1, \\ 10, & j_k = 2, \\ -10, & j_k = 3, \end{cases} \quad (46)$$

i.e., the random variable  $j_k$  randomly jumps; the diagonal elements of the transition probability matrix are set to 0.8s and the non-diagonal ones are set to 0.1s. Therefore, in state estimation for this jump linear system, at every time  $k$ , we need to jointly estimate the unknown state  $\mathbf{x}_k$  and the system's unknown operating mode  $j_k$  based on the past measurements  $\{y_1, y_2, \dots, y_k\}$ .

We reuse the real-world data and experimental setups from [37]: In short, data from usual GPS are to be processed to obtain higher-accuracy target positions and velocities in real time, while data from RTK serve as ground truth. The RTAMSE, for the total  $K$  time steps, is computed as

$$\sqrt{\frac{1}{K} \sum_{k=1}^K (p_k - \hat{p}_k)^2 + (s_k - \hat{s}_k)^2},$$

where  $p_k$  (resp.  $s_k$ ) denotes the true value of position (resp. velocity) at the time  $k$ , and  $\hat{p}_k$  (resp.  $\hat{s}_k$ ) denotes its estimate. Note that  $(p_k, s_k)$  for every  $k$  is provided by RTK. (Some authors prefer to independently report the RTAMSEs of position estimates  $\{\hat{p}_k\}_{\forall k \in [K]}$  and velocity estimates  $\{\hat{s}_k\}_{\forall k \in [K]}$ , respectively, because the two dimensions have different numerical scales. However, the author's experiments suggest that it introduces no essential influence on the main claims of this

paper. One may use the shared source codes to verify this point.)

**Track A Slowly-Maneuvering Car:** We first track a slowly-maneuvering car that travels on a road in Beijing, China. The car and its trajectory are shown in Fig. 9.

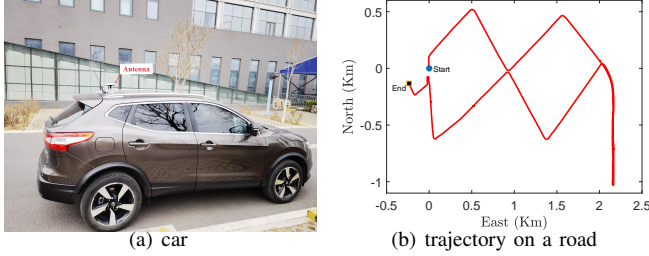


Fig. 9. Car and its trajectory (data credit: UniStrong Co., Ltd. and [37]).

The east axis (in the east-north-up coordinate) is investigated [37]. The RTAMSE performance against the values of  $(\alpha, \beta)$  is shown in Fig. 10.




Fig. 10. Averaged-RTAMSE performance of the  $(\alpha, \beta)$ -uncertainty-aware IMM filter, against the values of  $(\alpha, \beta)$ . The  $(\alpha, \beta)$ -uncertainty-aware IMM filter with  $(\alpha, \beta) = (2, 1.7)$  works best.

The results suggest that the  $(\alpha, \beta)$ -uncertainty-aware IMM (IMM-UA) filter with  $(\alpha, \beta) = (2, 1.7)$  works best, under which the tracking results (RTAMSE) are shown in Table I.

TABLE I  
TRACKING RESULTS OF THE CAR:  $(\alpha, \beta) = (2, 1.7)$

Filter	RTAMSE	Avg Time	Filter	RTAMSE	Avg Time
IMM	2.30	1.05e-05	IMM-UA	1.69	1.07e-05

Avg Time: Average Execution Time at each time step (unit: seconds).

From Table I, we can see that the value pair of  $(\alpha, \beta) = (2, 1.7)$  significantly reduces the tracking errors for the car. This value pair concentrates both the prior distribution (i.e., prior model weights) and the likelihood distribution (i.e., model likelihoods), which means that one of the three models in (46) dominates the rest two models most of the time. This implication coincides with our intuition from Fig. 9(b) that most of the time  $k$ , the model with  $a_{1,k} = 0$  (i.e., constant velocity and straight-line trajectory) dominates the motion of the car.

**Track A Highly-Maneuvering Drone:** We next track a highly-maneuvering drone that flies following round trajectories in the air over an open playground, with a flying speed

of about 6m/s during data collection. The drone and parts of its trajectory are shown in Fig. 11.




Fig. 11. Drone and its flying trajectory, starting from 10s and ending at 60s (data credit: Northwestern Polytechnical University and [37]).

The east axis (in the east-north-up coordinate) is investigated [37]. The RTAMSE performance against the values of  $(\alpha, \beta)$  is shown in Fig. 12.




Fig. 12. Averaged-RTAMSE performance of the  $(\alpha, \beta)$ -uncertainty-aware IMM filter, against the values of  $(\alpha, \beta)$ . The  $(\alpha, \beta)$ -uncertainty-aware IMM filter with  $(\alpha, \beta) = (0.2, 1.2)$  works best.

The results suggest that the  $(\alpha, \beta)$ -uncertainty-aware IMM filter with  $(\alpha, \beta) = (0.2, 1.2)$  works best, under which the tracking results (RTAMSE) are shown in Table II.

TABLE II  
TRACKING RESULTS OF THE DRONE:  $(\alpha, \beta) = (0.2, 1.2)$

Filter	RTAMSE	Avg Time	Filter	RTAMSE	Avg Time
IMM	0.77	2.71e-05	IMM-UA	0.60	2.85e-05

Avg Time: Average Execution Time at each time step (unit: seconds).

From Table II, we can see that the value pair of  $(\alpha, \beta) = (0.2, 1.2)$  significantly reduces the tracking errors for the drone. This value pair improves the entropy (i.e., the spread) of the likelihood distribution (i.e., the model likelihoods) and almost does not influence the prior distribution (i.e., the prior model weights because  $\beta = 1.2 \approx 1$ ), which means that none of the three models in (46) dominates the rest two models and the model set in (46) is not complete (viz., more candidate values for  $j_k$  and  $a_{j_k,k}$  are expected; only three values are not sufficient).<sup>6</sup> This implication coincides with our intuition from Fig. 11(b) that, since the drone highly maneuvers with acceleration quickly switching between positive values and

<sup>6</sup>For example,  $a_{j_k,k} := \{0, -2.5, -5, -7.5, -10, 2.5, 5, 7.5, 10\}$  should be better; however, this introduces much more computational loads in IMM filter. For extensive reading on this point, see [37, Subsec. VI.A.2.; Tab. III].



negative values, there is no model that dominates the motion of the drone.

4) *Remarks:* From our experiments, we found that the grid search method with the step size of 0.01 is an empirical golden rule. As for  $\tau$  such that  $(\alpha, \beta) \in [0, \tau]^2$ , experiments suggest that  $\tau \leq 3$  is practically sufficient.

#### IV. CONCLUSIONS

To combat the potential model misspecifications in prior distributions (i.e., prior belief) and/or data distributions (i.e., data evidence), this paper proposes to generalize the conventional Bayes' rule to the uncertainty-aware Bayes' rule. The uncertainty-aware Bayes' rule balances the relative importance of the prior distribution and the likelihood distribution by simply taking the exponentiation of the prior distribution and the likelihood distribution. We show that this exponentiation operator essentially adjusts the entropy (i.e., the concentration, the spread) of the prior distribution and the likelihood distribution. Therefore, with different exponents, the prior distribution and/or the likelihood distribution can be upweighted (i.e., by reducing the entropy) or downweighted (i.e., by inflating the entropy) in computing the posterior. Compared to the existing maximum entropy scheme, the uncertainty-aware Bayes' rule does not introduce much additional computational burden because the exponentiation operation is computationally light. In addition, compared to the existing maximum entropy scheme and  $\alpha$ -posterior scheme, the uncertainty-aware Bayes' rule is able to combat the conservativeness (i.e., upweight the useful distributional information) of the employed prior distributions and the likelihood distributions. Simulated and real-world applications further advocate that both opportunism and conservatism are potentially useful in practice. However, the optimal parameters  $(\alpha, \beta)$  cannot be theoretically specified because they depend on true prior distribution and data-generating distribution, which are unknown in practice. Therefore, the grid search and surrogate-optimization-based methods can be used to find the best values empirically.

Future papers following this work are expected to investigate the posterior consistency and asymptotic properties (e.g., asymptotic normality or non-normality) of the derived uncertainty-aware Bayes' rule, i.e., the  $(\alpha, \beta)$ -posterior; cf. [6], [15]. However, these problems are more mathematically statistical problems than theoretical signal processing problems, and therefore, we leave them to mathematical statisticians.

#### APPENDIX A PROOF OF LEMMA 1

*Proof:* According to the variational Bayesian interpretation of the posterior distribution  $p(\theta|\mathbf{y})$  in (1) (i.e., maximizing the evidence lower bound [38, p. 862]),  $p(\theta|\mathbf{y})$  solves

$$\min_{q(\theta)} \text{KL}[q(\theta) \| p(\theta)] + \mathbb{E}_{\theta \sim q(\theta)} [-\ln p(\mathbf{y}|\theta)].$$

Let  $C_{\mathbf{y}} := \int p(\mathbf{y}|\theta) d\theta$ . The above optimization problem is equivalent, in the sense of the same minimizer, to

$$\begin{aligned} \min_{q(\theta)} \text{KL}[q(\theta) \| p(\theta)] + \mathbb{E}_{\theta \sim q(\theta)} \left[ -\ln \frac{p(\mathbf{y}|\theta)}{C_{\mathbf{y}}} \right], \\ = \min_{q(\theta)} \text{KL}[q(\theta) \| p(\theta)] + \mathbb{E}_{\theta \sim q(\theta)} [-\ln l(\theta)], \\ = \min_{q(\theta)} \text{KL}[q(\theta) \| p(\theta)] + \text{KL}[q(\theta) \| l(\theta)] + \text{Ent } q(\theta). \end{aligned}$$

This completes the proof.  $\square$

#### APPENDIX B PROOF OF THEOREM 1

*Proof:* We have

$$\begin{aligned} \min_{q(\theta)} \alpha_1 \text{KL}[q(\theta) \| p(\theta)] + \alpha_2 \text{KL}[q(\theta) \| l(\theta)] + \alpha_3 \text{Ent } q(\theta) \\ = \min_{q(\theta)} \int q(\theta) \ln \frac{q^{\alpha_1 + \alpha_2}(\theta)}{p^{\alpha_1}(\theta) l^{\alpha_2}(\theta) q^{\alpha_3}(\theta)} d\theta \\ = \min_{q(\theta)} \int q(\theta) \ln \frac{q^{\alpha_1 + \alpha_2 - \alpha_3}(\theta)}{p^{\alpha_1}(\theta) l^{\alpha_2}(\theta)} d\theta. \end{aligned}$$

When  $\alpha_3 = \alpha_1 + \alpha_2$ , any distribution supported on  $\Theta^*$  solves the above optimization problem, where  $\Theta^*$  contains all maximizers of  $\ln[p^{\alpha_1}(\theta) l^{\alpha_2}(\theta)]$ .

When  $\alpha_3 < \alpha_1 + \alpha_2$ , the above optimization problem is equivalent to

$$(\alpha_1 + \alpha_2 - \alpha_3) \cdot \min_{q(\theta)} \int q(\theta) \ln \frac{q(\theta)}{p^{\frac{\alpha_1}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta) \cdot l^{\frac{\alpha_2}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta)} d\theta,$$

which is further equivalent, in the sense of the same minimizer, to

$$\min_{q(\theta)} \text{KL} \left[ q(\theta) \left\| \frac{p^{\frac{\alpha_1}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta) \cdot l^{\frac{\alpha_2}{\alpha_1 + \alpha_2 - \alpha_3}}(\theta)}{C} \right\| \right],$$

where  $C$  is the normalizer. This completes the proof.  $\square$

#### APPENDIX C PROOF OF THEOREM 2

*Proof:* Let  $C_{\alpha} := \int h^{\alpha}(\theta) d\theta$ . We have

$$\begin{aligned} \text{Ent } h^{(\alpha)}(\theta) &= \int \frac{h^{\alpha}(\theta)}{\int h^{\alpha}(\theta) d\theta} \cdot -\ln \frac{h^{\alpha}(\theta)}{\int h^{\alpha}(\theta) d\theta} d\theta \\ &= \frac{1}{C_{\alpha}} \int h^{\alpha}(\theta) \cdot [-\ln h^{\alpha}(\theta) + \ln C_{\alpha}] d\theta \\ &= \ln C_{\alpha} + \frac{\alpha}{C_{\alpha}} \int h^{\alpha}(\theta) \cdot -\ln h(\theta) d\theta. \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{d \text{Ent } h^{(\alpha)}(\theta)}{d\alpha} &= \frac{\alpha}{C_{\alpha}^2} \cdot \left\{ \left[ \int h^{\alpha}(\theta) \ln h(\theta) d\theta \right]^2 - \right. \\ &\quad \left. \int h^{\alpha}(\theta) d\theta \cdot \int h^{\alpha}(\theta) \ln h(\theta) \ln h(\theta) d\theta \right\}. \end{aligned}$$

Since  $h^\alpha(\boldsymbol{\theta}) \geq 0$  for every  $\alpha$  and  $\boldsymbol{\theta}$ , according to the Cauchy–Schwarz inequality, we have

$$\begin{aligned} & \int \left[ \sqrt{h^\alpha(\boldsymbol{\theta})} \right]^2 d\boldsymbol{\theta} \int \left[ \sqrt{h^\alpha(\boldsymbol{\theta}) \cdot \ln h(\boldsymbol{\theta}) \cdot \ln h(\boldsymbol{\theta})} \right]^2 d\boldsymbol{\theta} \\ & \geq \left[ \int \sqrt{h^\alpha(\boldsymbol{\theta}) \cdot h^\alpha(\boldsymbol{\theta}) \cdot \ln h(\boldsymbol{\theta}) \cdot \ln h(\boldsymbol{\theta})} d\boldsymbol{\theta} \right]^2 \\ & = \left[ \int h^\alpha(\boldsymbol{\theta}) \cdot |\ln h(\boldsymbol{\theta})| d\boldsymbol{\theta} \right]^2 \\ & \geq \left[ \int h^\alpha(\boldsymbol{y}) \cdot \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} \right]^2. \end{aligned}$$

As a result,

$$\frac{d \text{Ent } h^{(\alpha)}(\boldsymbol{\theta})}{d\alpha} \leq 0.$$

If  $\alpha \neq 0$ , the equality holds if and only if  $h(\boldsymbol{\theta})$  is a uniform distribution. This completes the proof.  $\square$

#### APPENDIX D PROOF OF THEOREM 3

Before proving Theorem 3, we prepare with the following lemma.

*Lemma 2:* If  $\alpha > 1$ , then

$$\int h^{(\alpha)}(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} > 0;$$

if  $\alpha < 1$ , then

$$\int h^{(\alpha)}(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} < 0.$$

*Proof of Lemma 2:* Let  $C_\alpha := \int h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta}$ . Jeffrey's divergence (NB: it is symmetric) between  $h(\boldsymbol{\theta})$  and  $h^{(\alpha)}(\boldsymbol{\theta})$  can be given as

$$\begin{aligned} & \text{KL} [h^{(\alpha)}(\boldsymbol{\theta}) \| h(\boldsymbol{\theta})] + \text{KL} [h(\boldsymbol{\theta}) \| h^{(\alpha)}(\boldsymbol{\theta})] \\ & = \int \frac{h^\alpha(\boldsymbol{\theta})}{C_\alpha} \ln \frac{h^\alpha(\boldsymbol{\theta})}{C_\alpha \cdot h(\boldsymbol{\theta})} d\boldsymbol{\theta} + \int h(\boldsymbol{\theta}) \ln \frac{C_\alpha \cdot h(\boldsymbol{\theta})}{h^\alpha(\boldsymbol{\theta})} d\boldsymbol{\theta} \\ & = \frac{1}{C_\alpha} \left[ \int h^\alpha(\boldsymbol{\theta}) \ln h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int h^\alpha(\boldsymbol{\theta}) \ln C_\alpha d\boldsymbol{\theta} - \right. \\ & \quad \left. \int h^\alpha(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} \right] + \left[ \int h(\boldsymbol{\theta}) \ln C_\alpha d\boldsymbol{\theta} + \right. \\ & \quad \left. \int h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int h(\boldsymbol{\theta}) \ln h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta} \right] \\ & = \frac{\alpha - 1}{C_\alpha} \int h^\alpha(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \ln C_\alpha + \ln C_\alpha + \\ & \quad (1 - \alpha) \int h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ & = (\alpha - 1) \left[ \int h^{(\alpha)}(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} \right] \\ & \geq 0. \end{aligned}$$

The equality holds if and only if  $\alpha = 1$ . This completes the proof.  $\square$

Now we proceed to the proof of Theorem 3.

*Proof of Theorem 3:* Let  $C_\alpha := \int h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta}$ . We have

$$\begin{aligned} & \text{KL} [h(\boldsymbol{\theta}) \| h^{(\alpha)}(\boldsymbol{\theta})] \\ & = \int h(\boldsymbol{\theta}) \ln \frac{h(\boldsymbol{\theta})}{h^{(\alpha)}(\boldsymbol{\theta})} d\boldsymbol{\theta} \\ & = \int h(\boldsymbol{\theta}) \ln \frac{C_\alpha \cdot h(\boldsymbol{\theta})}{h^\alpha(\boldsymbol{\theta})} d\boldsymbol{\theta} \\ & = \ln C_\alpha + (\alpha - 1) \text{Ent } h(\boldsymbol{\theta}). \end{aligned}$$

Therefore,

$$\begin{aligned} & \frac{d \text{KL} [h(\boldsymbol{\theta}) \| h^{(\alpha)}(\boldsymbol{\theta})]}{d\alpha} \\ & = \int h^{(\alpha)}(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} + \text{Ent } h(\boldsymbol{\theta}) \\ & = \int h^{(\alpha)}(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta}. \end{aligned}$$

Then, by Lemma 2, the monotonicity is immediate.

In addition,

$$\begin{aligned} & \frac{d^2 \text{KL} [h(\boldsymbol{\theta}) \| h^{(\alpha)}(\boldsymbol{\theta})]}{d^2 \alpha} \\ & = \frac{1}{C_\alpha^2} \left\{ \int h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta} \cdot \int h^\alpha(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \right. \\ & \quad \left. \left[ \int h^\alpha(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} \right]^2 \right\} \\ & \geq 0. \end{aligned}$$

The last inequality is due to the Cauchy–Schwarz inequality; the equality holds if and only if  $h(\boldsymbol{\theta})$  is a uniform distribution; see Appendix C. This completes the proof.  $\square$

#### APPENDIX E PROOF OF THEOREM 4

*Proof:* Let  $C_\alpha := \int h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta}$ . We have

$$\begin{aligned} f(\alpha) & := \text{KL} [h_0(\boldsymbol{\theta}) \| h(\boldsymbol{\theta})] - \text{KL} [h_0(\boldsymbol{\theta}) \| h^{(\alpha)}(\boldsymbol{\theta})] \\ & = \int h_0(\boldsymbol{\theta}) \ln \frac{h_0(\boldsymbol{\theta})}{h(\boldsymbol{\theta})} d\boldsymbol{\theta} - \int h_0(\boldsymbol{\theta}) \ln \frac{h_0(\boldsymbol{\theta}) C_\alpha}{h^\alpha(\boldsymbol{\theta})} d\boldsymbol{\theta} \\ & = \int h_0(\boldsymbol{\theta}) \ln \frac{h^\alpha(\boldsymbol{\theta})}{h(\boldsymbol{\theta}) C_\alpha} d\boldsymbol{\theta} \\ & = (\alpha - 1) \int h_0(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \ln \int h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta}. \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{df(\alpha)}{d\alpha} & = \int h_0(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \frac{\int h^\alpha(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta}}{\int h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta}} \\ & = \int h_0(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \int h^{(\alpha)}(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta}. \end{aligned}$$

As a result,

$$\left. \frac{df(\alpha)}{d\alpha} \right|_{\alpha=1} = \int [h_0(\boldsymbol{\theta}) - h(\boldsymbol{\theta})] \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta}.$$

The above display equals to zero if and only if  $h_0(\boldsymbol{\theta}) - h(\boldsymbol{\theta}) \equiv 0$  almost everywhere for  $\boldsymbol{\theta}$ , due to the arbitrariness of  $h_0$  and  $h$ . As assumed, we have  $h_0 \neq h$ , and therefore,  $\left. \frac{df(\alpha)}{d\alpha} \right|_{\alpha=1} \neq 0$ .

In addition,

$$\begin{aligned} & \frac{d^2 f(\alpha)}{d\alpha^2} \\ & = -\frac{1}{C_\alpha^2} \left\{ \int h^\alpha(\boldsymbol{\theta}) d\boldsymbol{\theta} \cdot \int h^\alpha(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} - \right. \\ & \quad \left. \left[ \int h^\alpha(\boldsymbol{\theta}) \ln h(\boldsymbol{\theta}) d\boldsymbol{\theta} \right]^2 \right\} \\ & \leq 0, \end{aligned}$$



where the last inequality is due to the Cauchy–Schwarz inequality; the equality holds if and only if  $h(\theta)$  is a uniform distribution; see Appendix C.

Since  $h(\theta)$  is not a uniform distribution, we have  $\frac{d^2 f(\alpha)}{d\alpha^2} < 0$  for all  $\alpha \geq 0$ . That is,  $f(\alpha)$  is strictly concave on  $[0, \infty)$ . In addition, we have  $f(1) = 0$ . Therefore, as long as  $\frac{df(\alpha)}{d\alpha}|_{\alpha=1} \neq 0$ , there exists  $\alpha \in [0, \infty)$  such that  $f(\alpha) > 0$ . Note that  $f(\alpha)$  is continuous in  $\alpha$ .

This completes the proof.  $\square$

## REFERENCES

- [1] R. van de Schoot, S. Depaoli, R. King, B. Kramer, K. Märtens, M. G. Tadesse, M. Vannucci, A. Gelman, D. Veen, J. Willemsen *et al.*, “Bayesian statistics and modelling,” *Nature Reviews Methods Primers*, vol. 1, no. 1, p. 1, 2021.
- [2] J. O. Berger, “An overview of robust Bayesian analysis,” *Test*, vol. 3, no. 1, pp. 5–124, 1994.
- [3] D. R. Insua and F. Ruggeri, *Robust Bayesian Analysis*. New York: Springer, 2000.
- [4] F. Ruggeri, D. R. Insua, and J. Martín, “Robust Bayesian analysis,” *Handbook of Statistics*, vol. 25, pp. 623–667, 2005.
- [5] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. Chapman and Hall/CRC, 2013.
- [6] M. A. Medina, J. L. M. Olea, C. Rush, and A. Velez, “On the robustness to misspecification of  $\alpha$ -posteriors and their variational approximations,” *The Journal of Machine Learning Research*, vol. 23, no. 1, pp. 6579–6629, 2022.
- [7] I. R. Petersen and A. V. Savkin, *Robust Kalman Filtering for Signals and Systems with Large Uncertainties*. Springer Science & Business Media, 1999.
- [8] A. M. Zoubir, V. Koivunen, E. Ollila, and M. Muma, *Robust Statistics for Signal Processing*. Cambridge University Press, 2018.
- [9] Y. S. Shmaliy, F. Lehmann, S. Zhao, and C. K. Ahn, “Comparing Robustness of the Kalman,  $H_\infty$ , and UFIR Filters,” *IEEE Transactions on Signal Processing*, vol. 66, no. 13, pp. 3447–3458, 2018.
- [10] S. Wang, “Distributionally Robust State Estimation,” Ph.D. dissertation, National University of Singapore, 2022.
- [11] P. J. Huber and E. M. Ronchetti, *Robust Statistics*, 2nd ed. John Wiley & Sons Hoboken, NJ, USA, 2009.
- [12] P. J. Huber, “Robust estimation of a location parameter,” *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964. [Online]. Available: <http://www.jstor.org/stable/2238020>
- [13] F. R. Hampel, “The influence curve and its role in robust estimation,” *Journal of the American Statistical Association*, vol. 69, no. 346, pp. 383–393, 1974.
- [14] S. Ghosal and A. Van der Vaart, *Fundamentals of Nonparametric Bayesian inference*. Cambridge University Press, 2017, vol. 44.
- [15] P. Alquier and J. Ridgway, “Concentration of tempered posteriors and of their variational approximations,” *The Annals of Statistics*, vol. 48, no. 3, pp. 1475–1497, 2020. [Online]. Available: <https://doi.org/10.1214/19-AOS1855>
- [16] S. Wang, “Distributionally robust state estimation for nonlinear systems,” *IEEE Transactions on Signal Processing*, vol. 70, pp. 4408–4423, 2022.
- [17] V. Bordinon, V. Matta, and A. H. Sayed, “Adaptive social learning,” *IEEE Transactions on Information Theory*, vol. 67, no. 9, pp. 6053–6081, 2021.
- [18] P. Hu, V. Bordinon, S. Vlaski, and A. H. Sayed, “Optimal aggregation strategies for social learning over graphs,” *IEEE Transactions on Information Theory*, 2023.
- [19] M. J. Rufo, J. Martín, and C. J. Pérez, “Log-Linear Pool to Combine Prior Distributions: A Suggestion for a Calibration-Based Approach,” *Bayesian Analysis*, vol. 7, no. 2, pp. 411 – 438, 2012. [Online]. Available: <https://doi.org/10.1214/12-BA714>
- [20] G. Koliander, Y. El-Laham, P. M. Djurić, and F. Hlawatsch, “Fusion of probability density functions,” *Proceedings of the IEEE*, vol. 110, no. 4, pp. 404–453, 2022.
- [21] Y. Shen and C. A. Shoemaker, “Global optimization for noisy expensive black-box multi-modal functions via radial basis function surrogate,” in *2020 Winter Simulation Conference (WSC)*. IEEE, 2020, pp. 3020–3031.
- [22] X. Wang, Y. Jin, S. Schmitt, and M. Olhofer, “Recent advances in Bayesian optimization,” *ACM Computing Surveys*, vol. 55, no. 13s, pp. 1–36, 2023.
- [23] R. Garnett, *Bayesian Optimization*. Cambridge University Press, 2023.
- [24] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, “Learning word vectors for sentiment analysis,” in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Portland, Oregon, USA: Association for Computational Linguistics, June 2011, pp. 142–150. [Online]. Available: <http://www.aclweb.org/anthology/P11-1015>
- [25] Scikit-learn. Naive Bayes (scikit-learn documentation). [Online]. Available: [https://scikit-learn.org/stable/modules/naive\\_bayes.html](https://scikit-learn.org/stable/modules/naive_bayes.html)
- [26] W. Wolberg, O. Mangasarian, N. Street, and W. Street, “Breast Cancer Wisconsin (Diagnostic),” UCI Machine Learning Repository, 1995, DOI: 10.24432/C5DW2B.
- [27] B. D. Anderson and J. B. Moore, *Optimal Filtering*. Prentice-Hall, N.J., 1979.
- [28] S. Wang, Z. Wu, and A. Lim, “Robust state estimation for linear systems under distributional uncertainty,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 5963–5978, 2021.
- [29] S. Wang and Z.-S. Ye, “Distributionally robust state estimation for linear systems subject to uncertainty and outlier,” *IEEE Transactions on Signal Processing*, vol. 70, pp. 452–467, 2021.
- [30] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, “A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking,” *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [31] B. P. Carlin, N. G. Polson, and D. S. Stoffer, “A monte carlo approach to nonnormal and nonlinear state-space modeling,” *Journal of the American Statistical Association*, vol. 87, no. 418, pp. 493–500, 1992.
- [32] H. A. Blom and Y. Bar-Shalom, “The interacting multiple model algorithm for systems with Markovian switching coefficients,” *IEEE Transactions on Automatic Control*, vol. 33, no. 8, pp. 780–783, 1988.
- [33] Y. Bar-Shalom, S. Challa, and H. A. Blom, “Imm estimator versus optimal estimator for hybrid systems,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 3, pp. 986–991, 2005.
- [34] X. R. Li and V. P. Jilkov, “Survey of maneuvering target tracking. part i. dynamic models,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333–1364, 2003.
- [35] V. P. Jilkov and X. R. Li, “Online Bayesian estimation of transition probabilities for markovian jump systems,” *IEEE Transactions on Signal Processing*, vol. 52, no. 6, pp. 1620–1630, 2004.
- [36] S. Zhao and F. Liu, “Recursive estimation for Markov jump linear systems with unknown transition probabilities: A compensation approach,” *Journal of the Franklin Institute*, vol. 353, no. 7, pp. 1494–1517, 2016.
- [37] S. Wang, “Distributionally robust state estimation for jump linear systems,” *IEEE Transactions on Signal Processing*, pp. 3835–3851, 2023.
- [38] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *Journal of the American Statistical Association*, vol. 112, no. 518, pp. 859–877, 2017.



**Dr. Shixiong Wang** received the B.Eng. degree in detection, guidance, and control technology, and the M.Eng. degree in systems and control engineering from the School of Electronics and Information, Northwestern Polytechnical University, China, in 2016 and 2018, respectively. He received his Ph.D. degree from the Department of Industrial Systems Engineering and Management, National University of Singapore, Singapore, in 2022.

He is currently a Postdoctoral Research Associate with the Intelligent Transmission and Processing Laboratory, Imperial College London, London, United Kingdom, from May 2023. He was a Postdoctoral Research Fellow with the Institute of Data Science, National University of Singapore, Singapore, from March 2022 to March 2023.

His research interest includes statistics and optimization theories with applications in signal processing (especially optimal estimation theory), machine learning (especially generalization error theory), and control technology.