
A Dimensionality Reduction Method for Finding Least Favorable Priors with a Focus on Bregman Divergence

Alex Dytso
New Jersey Institute
of Technology

Mario Goldenbaum
Bremen University of
Applied Sciences

H. Vincent Poor
Princeton University

Shlomo Shamai (Shitz)
Technion - Israel
Institute of Technology

Abstract

A common way of characterizing minimax estimators in point estimation is by moving the problem into the Bayesian estimation domain and finding a least favorable prior distribution. The Bayesian estimator induced by a least favorable prior, under mild conditions, is then known to be minimax. However, finding least favorable distributions can be challenging due to inherent optimization over the space of probability distributions, which is infinite-dimensional. This paper develops a dimensionality reduction method that allows us to move the optimization to a finite-dimensional setting with an explicit bound on the dimension. The benefit of this dimensionality reduction is that it permits the use of popular algorithms such as projected gradient ascent to find least favorable priors. Throughout the paper, in order to make progress on the problem, we restrict ourselves to Bayesian risks induced by a relatively large class of loss functions, namely Bregman divergences.

1 INTRODUCTION

Consider the problem of estimating a deterministic parameter $x \in \Omega \subseteq \mathbb{R}^n$ from a noisy observation $Y \in \mathcal{Y} \subseteq \mathbb{R}^k$, where x and Y are related through a conditional distribution $P_{Y|X=x}$. The standard objective in estimation theory is to find an estimator $f(Y)$ that minimizes some risk function $r(x, f)$. Formally, the risk function can be defined as

$$r(x, f) := \mathbb{E}[\ell(x, f(Y))], \quad (1)$$

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

where the expectation is taken with respect to $P_{Y|X=x}$ and where $\ell : \Omega \times \mathbb{R}^n \rightarrow [0, \infty)$ is some loss function (e.g., square loss). A common design principle is to look for an estimator that achieves the smallest maximum risk among all estimators. Such an estimator is called *minimax*. More precisely, an estimator $f_M : \mathcal{Y} \rightarrow \mathbb{R}^n$ is said to be minimax if

$$\sup_{x \in \Omega} r(x, f_M) = \inf_f \sup_{x \in \Omega} r(x, f), \quad (2)$$

where the infimum is taken over all measurable functions.

A potential way of finding the minimax estimator is to go along the Bayesian route. More specifically, consider the problem of estimating a random vector $X \in \Omega \subseteq \mathbb{R}^n$ with prior distribution P_X from a noisy observation $Y \in \mathcal{Y} \subseteq \mathbb{R}^k$ that are related through the same conditional distribution $P_{Y|X}$. The minimum Bayesian risk for prior P_X is then defined as

$$R(P_X, P_{Y|X}) := \inf_f \mathbb{E}[r(X, f)], \quad (3)$$

where the expectation is taken with respect to P_X . Furthermore, a prior distribution P_X^* is said to be *least favorable* if

$$R(P_X^*, P_{Y|X}) = \sup_{P_X} R(P_X, P_{Y|X}), \quad (4)$$

where the supremum is taken over all distributions supported on Ω . In other words, the random vector that follows P_X^* is the ‘hardest’ to estimate. A classical result in estimation theory states that

$$f_{P_X^*} \in \arg \inf_f \mathbb{E}_{P_X^*}[r(X, f)]; \quad (5)$$

that is, a best estimator for the least favorable prior is also a minimax estimator. See for instance [Lehmann and Casella, 1998] where this is shown under very mild conditions. Due to this fact, finding least favorable prior distributions has received considerable attention in the literature.

However, finding least favorable priors is a formidable task. The difficulty stems from the fact that the optimization in (4) is done over the space of probability distributions, which results in an infinite-dimensional optimization problem. The objective of this work is to show that under mild conditions the optimization problem in (4) can be reduced to a finite dimensional one. The key benefit of such a reduction is that one can begin to use numerical recipes to find a least favorable prior (e.g., gradient ascent algorithm).

The outline and the contributions of the paper are as follows. The remaining part of Section 1 is dedicated to notational remarks and past work. Section 2 presents some preliminary definitions (e.g., Bregman divergence), provides the problem statement and discusses our assumptions. Section 3 presents our main results, which show that under certain general conditions finding a least favorable prior can be reduced to a finite-dimensional optimization problem. Section 4 is devoted to the proof of the main theorem. Section 5 builds on the results of Section 3 and discusses how a projected gradient ascent algorithm can be used to find least favorable prior distributions. The algorithm is then applied to find least favorable priors in the context of binomial noise and quantized Gaussian noise.

1.1 Notational Remarks

Deterministic scalars and vectors are denoted by lower case letters and random objects by capital letters; \mathbb{R} denotes the affinely extended real number system; $\|\cdot\|$ denotes the Euclidean norm; the closed ball in \mathbb{R}^n of radius r centered at x is denoted as $\mathcal{B}_x(r) := \{y \in \mathbb{R}^n : \|y - x\| \leq r\}$; for a random vector $X \in \mathbb{R}^n$ and every measurable set $\mathcal{A} \subset \mathbb{R}^n$ we denote the probability measure of X as $P_X(\mathcal{A}) = \mathbb{P}[X \in \mathcal{A}]$; if it is clear from the context, we sometimes write P instead of P_X ; the space of all probability measures defined on sample space $\Omega \subseteq \mathbb{R}^n$ is denoted as $\mathcal{P}(\Omega)$; the Dirac measure centered on a fixed point x is denoted as δ_x ; for two probability distributions P and Q , $P \ll Q$ means P is absolutely continuous with respect to Q ; for a random vector X with distribution P_X the expected value is $\mathbb{E}[X] = \int x dP_X(x)$ and when we need to emphasize that X is distributed according to P_X we use $\mathbb{E}_{P_X}[X]$.

1.2 Past Work

The theory of finding least favorable prior distributions has received some attention for the special case when the noise is Gaussian and the loss function quadratic, for which $R(P_X, P_{Y|X})$ is commonly known as the *minimum mean square error*. For the univariate case (i.e., $n = 1$), Ghosh has shown in [Ghosh, 1964] that if the support of X is bounded

and the noise is Gaussian (i.e., $P_{Y|X} = \mathcal{N}(x, 1)$), then least favorable priors are discrete with finitely many mass points. Also for $P_{Y|X} = \mathcal{N}(x, 1)$, the authors of [Casella and Strawderman, 1981] capitalized on the result of Ghosh and provided necessary and sufficient conditions for the optimality of a two mass points prior distribution, and sufficient conditions for the optimality of three mass points priors. In [Berry, 1990], Berry has extended the results of [Casella and Strawderman, 1981] to the case of multivariate Gaussian noise with covariance taken to be the identity matrix. For $n \geq 1$, the authors of [Dytso et al., 2018b] have considered generalized moment constraints or linear constraints (i.e., $\mathbb{E}[g(X)] \leq c$) on X , and have shown that if $g \in o(\|x\|^2)$, then the support of a least favorable distribution is unbounded, and if $g \in \omega(\|x\|^2)$ it is bounded.

Much less work has been done for the general case. In [Kempthorne, 1987], for instance, it has been shown for $n = 1$ that if the conditional Bayes risk $\mathbb{E}[\ell(X, f(Y)) | X = x]$ is an analytic function of x and the support is bounded, then the least favorable prior is discrete with finitely many mass points. For a summary of known results on the properties of least favorable priors together with some extensions the interested reader is referred to [Marchand et al., 2004].

Algorithms for computing least favorable priors have been proposed in [Kempthorne, 1987, Nelson, 1966] and have been shown to converge under certain conditions. However, as these algorithms were designed without an explicit upper bound on the number of mass points in the support, their procedure relies on an optimization over an infinite dimensional space. Furthermore, it is not difficult to show that the cutting-plane algorithm, proposed in [Huang and Meyn, 2005] for finding the capacity-achieving distribution of a communication channel, can be adapted to the setting of finding a least favorable prior. However, it shares the same drawback as the algorithms proposed in [Kempthorne, 1987, Nelson, 1966] as it also relies on an optimization over an infinite dimensional space.

Finally, our method is inspired by the dimensionality reduction studied in the context of mutual information in [Witsenhausen, 1980], from which we borrow several key ideas (e.g., Dubins' theorem). See also [Dytso et al., 2018a].

2 PRELIMINARIES AND PROBLEM STATEMENT

2.1 Bregman Divergence and Bayesian Risk

To even have a chance to solve the optimization problem in (4), we need to slightly restrict the class of loss

functions. To that end, we will consider the following class, which is in fact very large.

Definition 1. (Bregman Divergence) *Let $\phi : \Omega \rightarrow \mathbb{R}$ be a continuously differentiable and strictly convex function. The Bregman divergence associated with ϕ is then defined as*

$$\ell_\phi(u, v) = \phi(u) - \phi(v) - \langle u - v, \nabla\phi(v) \rangle. \quad (6)$$

The classical squared error loss is recovered through Bregman divergences by choosing $\phi(u) = \|u\|^2$. As another example, consider the function $\phi(u) = u_1 \log u_1 + u_2 \log u_2$ where $u = [u_1, u_2]^\top$ with $\Omega = \mathbb{R}_+^2$, which induces the following Bregman divergence (known as the generalized I-divergence): for $u = [u_1, u_2]^\top$ and $v = [v_1, v_2]^\top$

$$\ell_\phi(u, v) = u_1 \log \frac{u_1}{v_1} + u_2 \log \frac{u_2}{v_2} - (u_1 - v_1) - (u_2 - v_2). \quad (7)$$

Bregman divergences have been introduced in [Bregman, 1967] in the context of convex optimization. In [Csiszár, 1991], Bregman divergences, together with f -divergences, were characterized axiomatically and considered in estimation settings. A thorough investigation of their properties was undertaken in [Banerjee et al., 2005b], where it was shown that many commonly used loss functions are members of this family. The structure of the optimal estimator under Bregman divergences as loss functions was studied in [Banerjee et al., 2005a], where it was shown that the conditional expectation is the unique minimizer. For extensions of Bregman divergence to different spaces the interested reader is referred to [Frigyik et al., 2008, Iyer and Bilmes, 2012, Wang et al., 2014] and references therein.

Definition 2. (Bayesian Risk with Respect to Bregman Divergence.) *For a joint distribution P_{XY} , we denote the Bayesian risk with respect to loss function ℓ_ϕ as*

$$R_\phi(P_X, P_{Y|X}) := \inf_{f: f \text{ is measurable}} \mathbb{E}[\ell_\phi(X, f(Y))]. \quad (8)$$

The following theorem summarizes some fundamental properties of ℓ_ϕ and $R_\phi(P_X, P_{Y|X})$, the proof of which can be found in [Banerjee et al., 2005b] and [Banerjee et al., 2005a].

Theorem 1. (Fundamental Properties of Bregman Divergence and Bayesian Risk)

1. (Non-Negativity) $\forall u, v \in \Omega : \ell_\phi(u, v) \geq 0$, with equality if and only if $u = v$;
2. (Convexity) $\ell_\phi(u, v)$ is convex in u ;

3. (Linearity) $\ell_\phi(u, v)$ is linear in ϕ ;
4. (Orthogonality Principle and Pythagorean Identity) For every random variable $X \in \Omega$ and every $u \in \Omega$

$$\mathbb{E}[\ell_\phi(X, u)] = \mathbb{E}[\ell_\phi(X, \mathbb{E}[X])] + \ell_\phi(\mathbb{E}[X], u). \quad (9)$$

Moreover, for any $f(Y)$

$$\mathbb{E}[\ell_\phi(X, f(Y))] = \mathbb{E}[\ell_\phi(X, \mathbb{E}[X|Y])] + \mathbb{E}[\ell_\phi(\mathbb{E}[X|Y], f(Y))]. \quad (10)$$

5. (Conditional Expectation is a Unique Bayesian Minimizer) If $\mathbb{E}[X] < \infty$ and $\mathbb{E}[\phi(X)] < \infty$, then

$$\begin{aligned} R_\phi(P_X, P_{Y|X}) &= \inf_{f: f \text{ is measurable}} \mathbb{E}[\ell_\phi(X, f(Y))] \\ &= \mathbb{E}[\ell_\phi(X, \mathbb{E}[X|Y])]. \end{aligned} \quad (11)$$

The optimizer in (11) is unique almost surely P_Y .

2.2 Moment Set

Definition 3. (Moment Set) *Let $(\Omega, \sigma(\Omega))$ be a measurable space and let $\mathcal{P}_{\text{reg}}(\Omega)$ be the set of all regular probability measures over the sample space Ω .¹ For any given $k \in \mathbb{N}$ fix measurable functions f_1, \dots, f_k as well as real numbers c_1, \dots, c_k . Then, the set*

$$\mathcal{H}_k := \{P \in \mathcal{P}_{\text{reg}}(\Omega) : \mathbb{E}_P[f_i(X)] \leq c_i, 1 \leq i \leq k\}; \quad (12)$$

that is, the set of regular probability measures with k bounded moments, is called the moment set.

Remark 1. *The restriction to regular probability measures is rather mild. For example, Ulam's theorem [Dudley, 2002, Th. 7.1.4] shows that a probability measure defined over a complete separable metric sample space (e.g., $\Omega = \mathbb{R}^n$) is regular.*

2.3 Conditional Expectation

Note that the conditional expectation $\mathbb{E}[X|Y]$ depends on the joint distribution P_{XY} through the conditional $P_{X|Y}$. However, since in this paper $P_{Y|X}$ is fixed and P_X varies, it is more convenient to treat $\mathbb{E}[X|Y]$ as a functional of P_X . Therefore, whenever we need to emphasize the dependence of the conditional expectation on the prior distribution, with a slight abuse of notation we will write $\mathbb{E}_{P_X}[X|Y]$.

Definition 4. *Consider a fixed $P_{Y|X}$ and a set of probability distributions $\mathcal{F} \subseteq \mathcal{P}(\Omega)$. We say that*

¹Recall that a probability measure is *regular* if any element of the σ -algebra $\sigma(\Omega)$ can be approximated from below by compact measurable sets and from above by open measurable sets.

Tweedie compatibility holds (with respect to $P_{Y|X}$ and \mathcal{F}), or T-compatibility for short, if there exists an operator $f : \mathcal{Y} \times \mathcal{P}(\Omega) \rightarrow \mathbb{R}^n$ such that for every $P_X \in \mathcal{F}$

$$\mathbb{E}_{P_X}[X|Y] = f(Y; P_Y) \text{ a.s.}, \quad (13)$$

where P_Y is the marginal distribution of Y induced by P_X (i.e., $\forall \mathcal{A} \in \sigma(\Omega) : P_Y(\mathcal{A}) = \mathbb{E}_{P_X}[P_{Y|X}(\mathcal{A}|X)]$).

T-compatibility simply says that the conditional expectation depends only on the marginal P_Y . An identity as in (13) is commonly known as Tweedie's formula [Robbins, 1956, Good, 1953]. A family of distributions that are T-compatible is the following.

Example 1. Consider an exponential family

$$P_{Y|\Theta}(y|\theta) = h(y)e^{y\theta - \psi(\theta)} \quad (14)$$

where θ is the natural parameter of the family, $h(y)$ the base measure, and $\psi(\theta)$ the log-partition function [Barndorff-Nielsen, 1978]. Now, if we set $X = e^\Theta$, then for $y \in \mathcal{Y}$

$$\mathbb{E}[X|Y = y] = \mathbb{E}[e^\Theta|Y = y] = \frac{h(y)}{h(y+1)} \frac{P_Y(y+1)}{P_Y(y)}. \quad (15)$$

For a concrete example consider $\Omega = [0, \infty)$, $\mathcal{Y} = \mathbb{N} \cup \{0\}$, and let $P_{Y|X}(y|x)$ be Poisson transition probabilities; that is,

$$P_{Y|X}(y|x) = \frac{1}{y!} x^y e^{-x}, \quad y \in \mathcal{Y}, \quad x \in \Omega, \quad (16)$$

where x is the mean parameter. It is not difficult to check that $\Theta = \log(X)$ and $h(y) = \frac{1}{y!}$ so that

$$\mathbb{E}[X|Y = y] = \frac{(y+1)P_Y(y+1)}{P_Y(y)}, \quad y \in \mathcal{Y}. \quad (17)$$

2.4 Problem Statement

We begin by listing assumptions that we are going to make throughout the rest of this paper.

Assumption 1.

- (i) \mathcal{H}_k as defined in (12) is compact;
- (ii) $R_\phi(P_X, P_{Y|X})$ is upper semicontinuous over \mathcal{H}_k ;
- (iii) $|\mathcal{Y}| \leq N$, $N \in \mathbb{N}$ (i.e., the support of the noisy observation Y is finite);
- (iv) For every $y \in \mathcal{Y}$, $P_{Y|X}(y|x)$ is continuous in x on the interior of $\Omega \subseteq \mathbb{R}^n$.

Note that assumptions (i), (ii) and (iv) are not very restrictive. For example, (i) and (ii) just guarantee the existence of a least favorable prior. Condition (iii), however, which imposes a restriction on the cardinality

of the support of noisy observation Y , represents the main restriction in this work. In many situations, \mathcal{Y} is indeed finite and in such case condition (iii) is not a limitation. See Section 5.2 for an example.

It is also important to emphasize that we do not impose any conditions on the positions taken by the support \mathcal{Y} of Y . Moreover, N can be taken as large as needed. Thus, Y can serve as an N -level quantization of some random vector U fully supported on \mathbb{R} . See Section 5.3 for an example.

Another assumption that we may or may not make is the following.

Assumption 2. T-compatibility, as defined in Definition 4, holds with respect to $P_{Y|X}$ and \mathcal{H}_k .

Now, the objective of this work is to study the optimization problem

$$\sup_{P_X \in \mathcal{H}_k} R_\phi(P_X, P_{Y|X}), \quad (18)$$

s.t. Assumption 1 and potentially also Assumption 2.

3 MAIN RESULT

The main result of this work is the following.

Theorem 2. (Least Favorable Distribution) *Let $N \in \mathbb{N}$ be finite, $k, n \in \mathbb{N}$ arbitrary but fixed, and Assumption 1 be fulfilled. Then, there exists a distribution $P_X^* \in \mathcal{H}_k$ with the following properties:*

- $\max_{P_X \in \mathcal{H}_k} R_\phi(P_X, P_{Y|X}) = R_\phi(P_X^*, P_{Y|X})$;
- $P_X^* \in \mathcal{H}_k$ is discrete with at most $N(k+1)(n+1)$ mass points (possibly containing mass points with individual coordinates equal to $\pm\infty$);
- $P_X^* \in \mathcal{H}_k$ is discrete with at most $N(k+1)$ mass points if in addition Assumption 2 is fulfilled.

Theorem 2 allows us to move the optimization in (4) from the space of probability distributions to \mathbb{R}^{nd+d} , where $d \leq N(k+1)(n+1)$ or $d \leq N(k+1)$ in case Assumption 2 is fulfilled. More specifically, we can parameterize the input distribution by a vector containing the sought after probability masses together with their locations:

$$P_X \Rightarrow \mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2]^\top \in \mathbb{R}^{nd+d} \\ = \left[\underbrace{x_1^\top, x_2^\top, \dots, x_d^\top}_{\substack{\text{points of the support} \\ x_i \in \mathbb{R}^n, 1 \leq i \leq d}}, \underbrace{p_1, p_2, \dots, p_d}_{\text{probability masses}} \right]^\top. \quad (19)$$

Working in \mathbb{R}^{nd+d} has the huge advantage that in order to find a least favorable distribution we can employ

numerical methods such as projected gradient ascent [Shalev-Shwartz and Ben-David, 2014]. More details on this will be given in Section 5.

Perhaps somewhat remarkable is that under Assumption 2 we obtain a bound on the cardinality of the support of a least favorable distribution that is *independent* of the dimension of Ω . This might have potential applications to the case where Ω is an infinite dimensional set.

We next show that under Assumption 1 and some extra conditions the dimensionality can further be reduced. The corresponding proof can be found in Section A of the supplementary material.

Proposition 1. *Let $N \in \mathbb{N}$ be finite, $k, n \in \mathbb{N}$ arbitrary but fixed, and Assumption 1 be fulfilled. Then, P_X^* is discrete with at most $(n+1)(N-1)+k+1$ mass points if*

- $\Omega \subset \mathbb{R}^n$ is compact and f_1, \dots, f_k are bounded and continuous on Ω ; or
- f_1, \dots, f_k are continuous on $\Omega = \mathbb{R}^n$ and are such that for every P_X with a finite number of mass points $\mathbb{E}_{P_X}[f_i(X)] < \infty$ implies $P_{X_j}(+\infty) = P_{X_j}(-\infty) = 0$ for all $1 \leq j \leq n$, with X_j the j -th coordinate of X .

Remark 2. *If functions f_1, \dots, f_k in the definition of moment set \mathcal{H}_k prevent the occurrence of mass points at $\pm\infty$, then the bound on the number of points can be reduced to $(n+1)(N-1)+k+1$. An example of such a function is $f : \mathbb{R} \rightarrow \mathbb{R}$, $f(x) = |x|^r$, $r > 0$, which naturally forces probability measures with a finite number of mass points to have mass points at $\pm\infty$ with zero probability.*

Finally, we note that Theorem 2 does not guarantee that the least favorable distribution is unique or that every least favorable distribution is discrete. It only guarantees that there exists a least favorable distribution that is discrete.

4 PROOF OF THE MAIN THEOREM

4.1 Preliminaries

Before actually proving Theorem 2, we provide some preliminary definitions and results that will help us to accomplish this.

4.1.1 Weak Convergence and Weak Continuity

It is well known that there exist several definitions of the convergence of a sequence of probability measures.

One is weak convergence, which provides a given space of probability measures with a topology.

Definition 5. *A sequence of probability measures $\{P_n\}_{n \in \mathbb{N}}$ is said to converge weakly to probability measure P if for every bounded and continuous function ψ*

$$\lim_{n \rightarrow \infty} \mathbb{E}_{P_n}[\psi(X)] \rightarrow \mathbb{E}_P[\psi(X)]. \quad (20)$$

Another main ingredient of our considerations are linear functionals. The following theorem gives a necessary and sufficient condition for a linear functional to be weakly continuous [Huber, 1981, Lemma 2.1].

Theorem 3. (Weak Continuity of Linear Functionals) *A linear functional $L : \mathcal{P}(\Omega) \rightarrow \mathbb{R}$ is weakly continuous on $\mathcal{P}(\Omega)$ if and only if there exists a bounded and continuous function ψ such that L can be represented as*

$$L(P) = \mathbb{E}_P[\psi(X)].$$

4.1.2 Linear Programming

The extreme value theorem for real-valued continuous functions over compact intervals is one of the most celebrated results of calculus. The following theorem is a generalization to compact topological spaces [Luenberger, 1997, Sec. 2.13].

Theorem 4. (Extreme Value Theorem) *For every compact topological space \mathcal{P} and every upper semicontinuous (lower semicontinuous) functional $f : \mathcal{P} \rightarrow \mathbb{R}$*

$$\sup_{P \in \mathcal{P}} f(P) = \max_{P \in \mathcal{P}} f(P) \quad \left(\inf_{P \in \mathcal{P}} f(P) = \min_{P \in \mathcal{P}} f(P) \right).$$

Moreover, if f is strictly concave (strictly convex) the maximizer (minimizer) is unique.

Definition 6. *An extreme point of any convex set \mathcal{S} is a point $x \in \mathcal{S}$ that cannot be represented as $x = (1-\alpha)y + \alpha z$ with $y, z \in \mathcal{S}$ and $\alpha \in (0, 1)$. We denote the set of all extreme points of \mathcal{S} as $\text{ex}(\mathcal{S})$.*

The following result states that when maximizing a linear functional over a moment set it is sufficient to focus on its extreme points [Winkler, 1988, Th. 3.2].

Theorem 5. (Linear Programming) *Let $L : \mathcal{P}(\Omega) \rightarrow \mathbb{R}$ be a linear functional. Then,*

$$\sup_{P \in \mathcal{H}_k} L(P) = \sup_{P \in \text{ex}(\mathcal{H}_k)} L(P).$$

Note that Theorem 5 only requires L to be linear and not necessarily continuous.

4.1.3 Extreme Points of a Moment Set

For proving the main result of this paper, the following theorem will be of particular importance [Winkler, 1988, Th. 2.1].

Theorem 6. (Extreme Points of a Moment Set) *For given $k \in \mathbb{N}$ let moment set \mathcal{H}_k be defined as in Definition 3. Then, the following holds:*

- \mathcal{H}_k is convex and the extreme points of \mathcal{H}_k are

$$\text{ex}(\mathcal{H}_k) \subseteq \overline{\text{ex}(\mathcal{H}_k)}, \quad (21)$$

where

$$\begin{aligned} & \overline{\text{ex}(\mathcal{H}_k)} \\ & := \left\{ P \in \mathcal{H}_k : P = \sum_{i=1}^m \alpha_i \delta_{x_i}, x_i \in \Omega, \alpha_i \in [0, 1], \right. \\ & \quad \left. \sum_{i=1}^m \alpha_i = 1, 1 \leq m \leq k+1 \wedge \text{the vectors} \right. \\ & \quad \left. [f_1(x_i), \dots, f_k(x_i), 1]^\top, 1 \leq i \leq m, \text{ are} \right. \\ & \quad \left. \text{linearly independent} \right\}; \end{aligned}$$

- If the moment conditions in (12) are fulfilled with equality, then (21) holds with equality.

Remark 3. *Theorem 6 is also valid in case we have a constraint on the support but no moment constraint (i.e., $k = 0$). For example, let $\Omega = \mathcal{B}_0(r)$ for some $r > 0$. Then,*

$$\text{ex}(\mathcal{H}_0) = \{P \in \mathcal{H}_0 : P = \delta_x, x \in \mathcal{B}_0(r)\}.$$

In case of a second moment constraint (i.e., $\mathcal{H}_1 = \{P \in \mathcal{P}_{\text{reg}}(\Omega) : \mathbb{E}_P[\|X\|^2] \leq c\}$, $c \in \mathbb{R}$) we have

$$\begin{aligned} & \text{ex}(\mathcal{H}_1) \subseteq \overline{\text{ex}(\mathcal{H}_1)} \\ & = \{P \in \mathcal{H}_1 : P = (1 - \alpha)\delta_{x_1} + \alpha\delta_{x_2}, \alpha \in [0, 1], \\ & \quad x_1, x_2 \in \Omega, \|x_1\| \neq \|x_2\|\}. \end{aligned}$$

4.1.4 Dubins' Theorem

Definition 7. *A convex set \mathcal{S} of a vector space \mathcal{V} is called linearly closed (linearly bounded) if every straight line intersects with \mathcal{S} on a closed (bounded) subset of that line.*

With this definition in hand, a powerful theorem proven by Dubins [Dubins, 1962] is the following.

Theorem 7. (Dubins' Theorem) *Let $f : \mathcal{V} \rightarrow \mathbb{R}$ be a linear functional over a vector space \mathcal{V} and let*

$$\mathcal{L} = \{v \in \mathcal{V} : f(v) = c\},$$

for some $c \in \mathbb{R}$, be a hyperplane formed by f . Moreover, let \mathcal{I} be the intersection of a linearly closed and linearly bounded convex set $\mathcal{K} \subset \mathcal{V}$ with n hyperplanes. Then, every extreme point of \mathcal{I} is a convex combination of at most $n + 1$ extreme points of \mathcal{K} .

A remarkable property of Theorem 7 is that it also holds for the infinite dimensional case. See Fig. 1 for a finite-dimensional example of Dubins' Theorem.




Figure 1: Example of Dubins' Theorem for $n = 1$. Cube $\mathcal{K} \subset \mathbb{R}^3$, which is a linearly closed and linearly bounded convex set, intersects with a hyperplane in a triangle \mathcal{I} . An extreme point/vertex of \mathcal{I} (e.g., point A) belongs to an edge of \mathcal{K} . Thus, A is the convex combination of $n + 1 = 2$ extreme points of \mathcal{K} ; that is, a convex combination of points B and C .

4.2 Proof of Theorem 2: Existence of a Solution

Recall that the first part of the main result states that with Assumption 1 a least favorable prior exists. Observe that this follows from Theorem 4 due to the semicontinuity of Bayesian risk $R_\phi(P_X, P_{Y|X})$ and the compactness of \mathcal{H}_k . Hence, the supremum in (18) is attained by some $P_X^* \in \mathcal{H}_k$.

4.3 Proof of Theorem 2: Assumptions 1 & 2

We first proof the second part of Theorem 2 where Assumptions 1 and 2 are both fulfilled as it is easier.

As a least favorable prior, P_X^* , exists (not necessarily unique), let $P_Y(\cdot; \star)$ and $\mathbb{E}_\star[X|Y] : \mathcal{Y} \rightarrow \Omega$ denote the marginal of Y and the conditional expectation of X given Y induced by P_X^* .² Moreover, we define

$$\begin{aligned} \mathcal{P}^\star & := \{P_X \in \mathcal{H}_k : P_Y(y; \star) = P_Y(y; P_X), \forall y \in \mathcal{Y}\} \\ & = \{P_X \in \mathcal{H}_k : c_y^\star = P_Y(y; P_X), \forall y \in \mathcal{Y}\} \quad (22) \end{aligned}$$

as the set of least favorable priors that induce $P_Y(y; \star)$, where for ease of notation $c_y^\star := P_Y(y; \star)$.

Next, note that as a consequence of Assumption 2, if $P_X, Q_X \in \mathcal{P}^\star$, then for every $y \in \mathcal{Y}$

$$\mathbb{E}_{P_X}[X|Y=y] = \mathbb{E}_{Q_X}[X|Y=y] = \mathbb{E}_\star[X|Y=y]. \quad (23)$$

Furthermore, observe that \mathcal{P}^\star is the intersection of \mathcal{H}_k with $N - 1$ hyperplanes of the form

$$\mathcal{L}_i = \{P_X : c_{y_i} = \mathbb{E}_{P_X}[P_{Y|X}(y_i|X)]\}, 1 \leq i \leq N - 1, \quad (24)$$

²More precisely, for given P_X^* , $P_Y(\cdot; \star)$ and $\mathbb{E}_\star[X|Y]$ are just shorthand for $P_Y(\cdot; P_X^*)$ and $\mathbb{E}_{P_X^*}[X|Y]$.

where we have used that $P_Y(y; P_X) = \mathbb{E}_{P_X}[P_{Y|X}(y|X)]$, $y \in \mathcal{Y}$. Note that we omitted hyperplane \mathcal{L}_N in (24) as in the space of probability distributions everything sums up to one so that \mathcal{L}_N is redundant. Note also that each \mathcal{L}_i is a closed set, which follows from Theorem 3 and the fact that $P_{Y|X}(y_i|x)$ is bounded and continuous in x for each y_i (i.e., sets defined by continuous functions are closed).

Next, observe that

$$\begin{aligned} & \max_{P_X \in \mathcal{H}_k} \mathbb{E}_{P_X}[\ell_\phi(X, \mathbb{E}_{P_X}[X|Y])] \\ & \stackrel{a)}{=} \max_{P_X \in \mathcal{P}^*} \mathbb{E}_{P_X}[\ell_\phi(X, \mathbb{E}_{P_X}[X|Y])] \end{aligned} \quad (25)$$

$$\stackrel{b)}{=} \max_{P_X \in \mathcal{P}^*} \mathbb{E}_{P_X}[\ell_\phi(X, \mathbb{E}_*[X|Y])] \quad (26)$$

$$\stackrel{c)}{=} \max_{P_X \in \text{ex}(\mathcal{P}^*)} \mathbb{E}_{P_X}[\ell_\phi(X, \mathbb{E}_*[X|Y])], \quad (27)$$

where *a*) follows from using the fact that \mathcal{P}^* contains a least favorable prior, *b*) from the fact that $\mathbb{E}_{P_X}[X|Y] = \mathbb{E}_*[X|Y]$ for every $P_X \in \mathcal{P}^*$, and *c*) from Theorem 5 by observing that $P_X \mapsto \mathbb{E}_{P_X}[\ell_\phi(X, \mathbb{E}_*[X|Y])]$ is linear over \mathcal{P}^* and therefore $\mathbb{E}_*[X|Y]$ does not change on \mathcal{P}^* .

Finally, recall that \mathcal{P}^* consists of the intersection of $N - 1$ hyperplanes, defined in (24), with \mathcal{H}_k . Thus, as \mathcal{P}^* is a subset of \mathcal{H}_k , it follows from Theorem 7 that every extreme point of \mathcal{P}^* (or every point of $\text{ex}(\mathcal{P}^*)$) can be represented by a convex combination of at most $(N - 1) + 1 = N$ extreme points of \mathcal{H}_k . Due to Theorem 6, however, the extreme points are discrete distributions with at most $k + 1$ mass points, so that P_X^* consists of at most $N(k + 1)$ mass points.

4.4 Proof of Theorem 2: Assumption 1 only

The key to the proof of Theorem 2 under Assumptions 1 and 2 is the construction of a set \mathcal{P}^* over which the conditional expectation does not change and which can be written as an intersection of finitely many hyperplanes with a moment set. Assumption 2 allows us to do this without the dependence on the dimension n of Ω . In the general case (i.e., without Assumption 2), however, this does not seem to be possible, which is why we have to construct \mathcal{P}^* differently.

Towards this end, let $\mathbb{E}_*[X|Y] : \mathcal{Y} \rightarrow \Omega$ as in the previous subsection be induced by some least favorable prior P_X^* . Furthermore, note that the j 's component of $\mathbb{E}_*[X|Y]$ can be written as

$$\mathbb{E}_*[X_j|Y = y] = \frac{\mathbb{E}_*[X_j P_{Y|X_j}(y|X_j)]}{P_Y(y; \star)} \quad (28)$$

$$= \frac{\mathbb{E}_*[X_j P_{Y|X_j}(y|X_j)]}{\mathbb{E}_*[P_{Y|X}(y|X)]} \quad (29)$$

for every $y \in \mathcal{Y}$ and $j = 1, \dots, n$. The expression in (29) implies that the conditional expectation is a ratio of two linear functionals. Therefore, the set \mathcal{P}^* can be constructed as follows:

$$\begin{aligned} \mathcal{P}^* &= \{P_X \in \mathcal{H}_k : c_{ji}^* = \mathbb{E}_{P_X}[X_j P_{Y|X_j}(y_i|X_j)] \wedge \\ & \quad c_i^* = P_Y(y_i; P_X) \ 1 \leq j \leq n, \ 1 \leq i \leq N\}, \end{aligned} \quad (30)$$

where for ease of notation

$$c_i^* := P_Y(y_i; \star), \ 1 \leq i \leq N - 1, \quad (31)$$

$$c_{ji}^* := \mathbb{E}_*[X_j P_{Y|X_j}(y_i|X_j)], \ 1 \leq j \leq n, \ 1 \leq i \leq N. \quad (32)$$

Again, we omitted the hyperplane c_N^* in (31) due to the same reasons as mentioned in the context of (24).

By construction $\mathbb{E}_*[X|Y = y_i] = \mathbb{E}_{P_X}[X|Y = y_i]$ for every $P_X \in \mathcal{P}^*$. In addition, the set \mathcal{P}^* is the intersection of the $(n + 1)N - 1$ hyperplanes

$$\mathcal{L}_i^{(a)} = \{P_X : c_i^* = \mathbb{E}_{P_X}[P_{Y|X}(y_i|X)]\} \quad (33)$$

for $1 \leq i \leq N - 1$ and

$$\mathcal{L}_{ij}^{(b)} = \{P_X : c_{ji}^* = \mathbb{E}_{P_X}[X_j P_{Y|X_j}(y_i|X_j)]\} \quad (34)$$

for $1 \leq i \leq N$ and $1 \leq j \leq n$, respectively. Now, at this point, following the same line of arguments as in Section 4.3 we arrive at the conclusion that P_X^* consists of at most $(k + 1)(n + 1)N$ mass points.

5 NUMERICAL EXAMPLES

To demonstrate the findings of this paper, in this section we present two numerical examples carried out with the projected gradient ascent method. In the following subsection, we first provide a description of that method where for the sake of simplicity we only focus on the case of a support constraint. That is, Ω is bounded and moment constraints are not present (i.e., $k = 0$). As a result, the first bound on the cardinality becomes $N(n + 1)$ and the second bound just N .

5.1 Projected Gradient Ascent

Firm upper bounds on the number of mass points, such as the one given in Theorem 2, allow us to carry out the optimization over the space \mathbb{R}^{nd+d} , where d denotes the number of mass points, instead over the space of probability distributions. As mentioned in the discussion after Theorem 2, working in \mathbb{R}^{nd+d} has the advantage that we can employ numerical methods such as projected gradient ascent [Shalev-Shwartz and Ben-David, 2014]. A quick sketch of how to use projected gradient ascent for finding a least favorable prior follows next.

As described in (19), we parameterize prior distribution P_X by a vector $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2]^\top \in \mathbb{R}^{nd+d}$ containing the sought after probability masses together with their locations. From Theorem 2 we know that d can be at most $N(n+1)$, or N if in addition Assumption 2 holds. Now, let

$$g(\mathbf{x}) := R_\phi(P_X, P_{Y|X}) \quad (35)$$

and define

$$\mathcal{G} := \{\mathbf{x} \in \mathbb{R}^{nd+d} : \mathbf{x}_1 \in \Omega, \mathbf{x}_2 \in \mathcal{S}\} \quad (36)$$

as the constraint set with \mathcal{S} denoting the probability simplex. In view of Theorem 2, we then have

$$\max_{P_X: X \in \Omega} R_\phi(P_X, P_{Y|X}) = \max_{\mathbf{x} \in \mathcal{G}} g(\mathbf{x}). \quad (37)$$

Starting from an initial point $\mathbf{x}^{(1)} \in \mathcal{G}$, projected gradient ascent iterates the following equation until a stopping criterion is met:

$$\mathbf{x}^{(t+1)} = \text{proj}_{\mathcal{G}}(\mathbf{x}^{(t)} + \lambda \nabla g(\mathbf{x}^{(t)})), t \in \mathbb{N}, \quad (38)$$

where $\lambda > 0$ is some step size, $\nabla g(\mathbf{x})$ the gradient of g , and $\text{proj}_{\mathcal{G}}(\cdot)$ the projection operator that tries to find a point $\mathbf{x}^{(t+1)} \in \mathcal{G}$ that is closest to $\mathbf{x}^{(t)} + \lambda \nabla g(\mathbf{x}^{(t)})$ in squared Euclidean distance.

If Ω is convex, the projection of \mathbf{x}_1 onto Ω can be done efficiently by using the alternating projection method [Bauschke and Borwein, 1996], whereas an efficient implementation of the projection of \mathbf{x}_2 onto the probability simplex with complexity $\mathcal{O}(d \log d)$ can be found in [Wang and Carreira-Perpinán, 2013].

In order to implement the projected gradient ascent algorithm described above, we obviously have to compute the gradient of g . The following result provides the gradient for the special case of $n = 1$ and ϕ being the squared error loss. The corresponding proof can be found in Section B of the supplementary material.

Proposition 2. *Let $\phi : \Omega \rightarrow \mathbb{R}$, $\phi(x) = x^2$. Then, if g is differentiable we have for $i = 1, \dots, N$*


$$\frac{\partial}{\partial p_i} g(\mathbf{x}) = \mathbb{E}[(x_i - \mathbb{E}[X|Y])^2 | X = x_i], \quad (39)$$

$$\begin{aligned} \frac{\partial}{\partial x_i} g(\mathbf{x}) &= 2p_i(x_i - \mathbb{E}[\mathbb{E}[X|Y] | X = x_i]) \\ &+ p_i \sum_{j=1}^N P'_{Y|X}(y_j | x_i) \left((\mathbb{E}[X|Y = y_j])^2 \right. \\ &\left. - 2x_i \mathbb{E}[X|Y = y_j] \right). \end{aligned} \quad (40)$$


5.2 Binomial Model

Consider conditional probability mass function

$$P_{Y|X}(y|x) = \binom{m}{y} x^y (1-x)^{m-y}, x \in \Omega, y \in \mathcal{Y}. \quad (41)$$



(a) Support of least favorable prior P_X^* for different values of m .




(b) Least favorable priors for $m = 3, 6, 10$.

Figure 2: Least favorable priors for $P_{Y|X}$ being binomial with parameter m .

Here, x is the success probability, $\Omega = [0, 1]$, and $\mathcal{Y} = \{0, 1, \dots, m\}$, $m \in \mathbb{N}$. We treat parameter m of this distribution as known so that $N = m + 1$. This binomial model is very popular and has a wide range of applications (see [Neter et al., 1988] for details).

We are now interested in estimating the success probability, which we model as a random variable $X \in \Omega$ with distribution P_X . To compute least favorable prior P_X^* , we use the gradient ascent algorithm described in the previous subsection.

Fig. 2 depicts the outcome of the algorithm for m ranging from 1 to 10. More specifically, Fig. 2a shows the support of P_X^* as a function of m . It is interesting to note that the support is uniformly spaced but P_X^* is not the uniform distribution, which can be seen from Fig. 2b for $m = 3, 6, 10$.


 Figure 3: Quantization function Q_N for $N = 1, 2, 3$.

5.3 Quantized Gaussian

As another example consider the scenario in which we seek to estimate a random variable X that is embedded in Gaussian noise:

$$W = X + Z, \quad Z \sim \mathcal{N}(0, 1). \quad (42)$$

Suppose that instead of observing W directly, we only have access to a quantized version of W , which we denote as Y . Specifically, consider uniform quantization with clipping; that is, for a fixed integer N

$$y = Q_N(w) = \begin{cases} \lfloor w \rfloor & |w| < N \\ \text{sign}(w)N & \text{else} \end{cases}, \quad (43)$$

where $\lfloor \cdot \rfloor$ rounds to the nearest integer. Examples of Q_N are depicted in Fig. 3.

Note that in this scenario, Y is supported on $\mathcal{Y} = \{-N, \dots, 0, \dots, N\}$. Moreover, the conditional distribution of Y given X is for $N \geq 1$ given by

$$P_{Y|X}(y|x) = \begin{cases} \Phi(-N - x + \frac{1}{2}), & y = -N \\ \Phi(y - x + \frac{1}{2}) - \Phi(-x - \frac{1}{2}), & |y| \leq N - 1, \\ \Phi(-N + x + \frac{1}{2}), & y = N \end{cases} \quad (44)$$

where Φ denotes the cumulative distribution function of a standard Gaussian. Using Theorem 2, we have that the cardinality of the support of P_X^* is bounded by $2N + 1$.

Now, for this model we would like to find the least favorable prior under the assumption that $X \in \Omega = [-A, A]$, $A > 0$. The corresponding numerical results are depicted in Fig. 4 for $A = 5$ and different N .




 (a) Support of least favorable prior P_X^* for $A = 5$ and different values of N .

 (b) Least favorable priors for $A = 5$ and $N = 1, 3, 4$.

 Figure 4: Least favorable priors for $P_{Y|X}$ given in (44) as a result of quantization.

6 CONCLUSION

This work has developed a dimensionality reduction method for finding least favorable prior distributions. The method produces an explicit bound on the size of the support of a least favorable prior and thus transforms the corresponding optimization problem from an infinite-dimensional to a finite-dimensional one. The numerical benefits of this method have been demonstrated via several examples based on a projected gradient ascent algorithm. Finally, while the focus was on loss functions that belong to the family of Bregman divergences, for which the optimal estimator is the conditional mean, the method can be generalized to other loss functions that have a concrete optimal estimator. For example, the results can be extended to the L_1 case, where the optimal Bayesian estimator is also the conditional median.

Acknowledgments

We would like to thank the anonymous reviewers for their comments and suggestions, which helped to improve the quality of this paper. The work of H. V. Poor was supported by the U.S. National Science Foundation (NSF) within the Israel-US Binational program under grant CCF-1908308 and the work of S. Shamai (Shitz) by the US-Israel Binational Science Foundation (BSF) under grant BSF-2018710, respectively.

References

- [Banerjee et al., 2005a] Banerjee, A., Guo, X., and Wang, H. (2005a). On the optimality of conditional expectation as a Bregman predictor. *IEEE Trans. Inf. Theory*, 51(7):2664–2669.
- [Banerjee et al., 2005b] Banerjee, A., Merugu, S., Dhillon, I. S., Ghosh, J., and Lafferty, J. (2005b). Clustering with Bregman divergences. *J. Mach. Learn. Res.*, 6(10).
- [Barndorff-Nielsen, 1978] Barndorff-Nielsen, O. (1978). *Information and Exponential Families*. John Wiley & Sons.
- [Bauschke and Borwein, 1996] Bauschke, H. H. and Borwein, J. M. (1996). On projection algorithms for solving convex feasibility problems. *SIAM Rev.*, 38(3):367–426.
- [Berry, 1990] Berry, J. C. (1990). Minimax estimation of a bounded normal mean vector. *J. Multivariate Anal.*, 35(1):130–139.
- [Bregman, 1967] Bregman, L. M. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. & Math. Phys.*, 7(3):200–217.
- [Casella and Strawderman, 1981] Casella, G. and Strawderman, W. E. (1981). Estimating a bounded normal mean. *Ann. Statist.*, 9(4):870–878.
- [Csiszár, 1991] Csiszár, I. (1991). Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems. *Ann. Statist.*, 19(4):2032–2066.
- [Dubins, 1962] Dubins, L. E. (1962). On extreme points of convex sets. *Math. Anal. Appl.*, 5(2):237–244.
- [Dudley, 2002] Dudley, R. M. (2002). *Real Analysis and Probability*. Cambridge University Press.
- [Dytso et al., 2018a] Dytso, A., Goldenbaum, M., Poor, H. V., and Shamai (Shitz), S. (2018a). When are discrete channel inputs optimal? – Optimization techniques and some new results. In *Proc. 52nd Annu. Conf. Inf. Sci. Syst. (CISS)*, pages 1–6.
- [Dytso et al., 2018b] Dytso, A., Poor, H. V., Bustin, R., and Shamai, S. (2018b). On the structure of the least favorable prior distributions. In *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, pages 1081–1085.
- [Frigyik et al., 2008] Frigyik, B. A., Srivastava, S., and Gupta, M. R. (2008). Functional Bregman divergence and Bayesian estimation of distributions. *IEEE Trans. Inf. Theory*, 54(11):5130–5139.
- [Ghosh, 1964] Ghosh, M. N. (1964). Uniform approximation of minimax point estimates. *Ann. Math. Statist.*, 35(3):1031–1047.
- [Good, 1953] Good, I. J. (1953). The population frequencies of species and the estimation of population parameters. *Biometrika*, 40(3-4):237–264.
- [Huang and Meyn, 2005] Huang, J. and Meyn, S. P. (2005). Characterization and computation of optimal distributions for channel coding. *IEEE Trans. Inf. Theory*, 51(7):2336–2351.
- [Huber, 1981] Huber, P. J. (1981). *Robust Statistics*. John Wiley & Sons.
- [Iyer and Bilmes, 2012] Iyer, R. and Bilmes, J. A. (2012). Submodular-Bregman and the Lovász-Bregman divergences with applications. In *Adv. Neural Inf. Process Syst.*, pages 2933–2941.
- [Kempthorne, 1987] Kempthorne, P. J. (1987). Numerical specification of discrete least favorable prior distributions. *SIAM J. Sci. and Stat. Comp.*, 8(2):171–184.
- [Lehmann and Casella, 1998] Lehmann, E. L. and Casella, G. (1998). *Theory of Point Estimation*. Springer, 2nd edition.
- [Luenberger, 1997] Luenberger, D. G. (1997). *Optimization by Vector Space Methods*. John Wiley & Sons.
- [Marchand et al., 2004] Marchand, E., Strawderman, W. E., et al. (2004). Estimation in restricted parameter spaces: A review. *A Festschrift for Herman Rubin, IMS Lecture Notes Monogr. Ser.*, 45:21–44.
- [Nelson, 1966] Nelson, W. (1966). Minimax solution of statistical decision problems by iteration. *Ann. Math. Statist.*, 37(6):1643–1657.

- [Neter et al., 1988] Neter, J., Wasserman, W., and Whitmore, G. A. (1988). *Applied Statistics*. Allyn & Bacon.
- [Robbins, 1956] Robbins, H. (1956). An empirical Bayes approach to statistics. In *Proc. Berkeley Symp. Math. Statist. Probab.*, pages 157–163.
- [Shalev-Shwartz and Ben-David, 2014] Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- [Wang et al., 2014] Wang, L., Carlson, D. E., Rodrigues, M. R., Calderbank, R., and Carin, L. (2014). A Bregman matrix and the gradient of mutual information for vector Poisson and Gaussian channels. *IEEE Trans. Inf. Theory*, 60(5):2611–2629.
- [Wang and Carreira-Perpinán, 2013] Wang, W. and Carreira-Perpinán, M. A. (2013). Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. *arXiv preprint arXiv:1309.1541*.
- [Winkler, 1988] Winkler, G. (1988). Extreme points of moment sets. *Math. Oper. Res.*, 13(4):581–587.
- [Witsenhausen, 1980] Witsenhausen, H. S. (1980). Some aspects of convexity useful in information theory. *IEEE Trans. Inf. Theory*, 26(3):265–271.

Supplementary Material: A Dimensionality Reduction Method for Finding Least Favorable Priors with a Focus on Bregman Divergence

A PROOF OF PROPOSITION 1

Now that we know from Theorem 2 that the least favorable distribution is discrete with at most $(k+1)(n+1)N$ mass points, we are able to slightly refine the number of mass points under various additional conditions.

A.1 Refinement for Compact Ω

Suppose that Ω is a proper compact subset of \mathbb{R}^n . Moreover, as the optimal input distribution exists, there exist numbers $c'_m \leq c_m$ such that

$$\mathbb{E}_{P_X^*} [f_i(X)] = c'_m, \quad m = 1, \dots, k.$$

Thus, the definition of the set in (30) can be modified to

$$\mathcal{P}^* = \left\{ P_X \in \mathcal{P}(\Omega) : \mathbb{E}_{P_X} [X_j P_{Y|X_j}(y_i|X_j)] = c_{ji}^*, P_Y(y_i; P_X) = P_Y(y_i; \star), \mathbb{E}_{P_X} [f_i(X)] = c'_m \right. \\ \left. m = 1, \dots, k, j = 1, \dots, n, i = 1, \dots, N \right\}.$$

Clearly, $P_X^* \in \mathcal{P}^*$. Note that \mathcal{P}^* is an intersection of the hyperplanes defined in (33) and (34) with the hyperplanes

$$\mathcal{V}_i := \{ P_X \in \mathcal{P}(\Omega) : \mathbb{E}_{P_X} [f_i(X)] = c'_m \}, \quad i = 1, \dots, k.$$

Moreover, each \mathcal{V}_i is a closed set as f_i is bounded and continuous on Ω for each $i = 1, \dots, k$ (see Theorem 3). Therefore, applying Dubins' theorem to the set \mathcal{P}^* , it follows that the extreme points of \mathcal{P}^* can be represented as convex combinations of at most $(n+1)(N-1) + k + 1$ extreme points of $\mathcal{P}(\Omega)$. As the extreme points of $\mathcal{P}(\Omega)$ are point masses, we have that the least favorable distribution P_X^* has at most $(n+1)(N-1) + k + 1$ mass points.

A.2 Refinement for Functions that Imply $P_{X_j}(+\infty) = P_{X_j}(-\infty) = 0$ for all $1 \leq j \leq n$

Now, let Ω be arbitrary and f_1, \dots, f_k bounded and continuous on Ω such that for every P_X with a finite number of mass points the conditions $\mathbb{E}_{P_X} [f_i(X)] < \infty$, $i = 1, \dots, k$, imply that $P_{X_j}(+\infty) = P_{X_j}(-\infty) = 0$ for all $j = 1, \dots, n$. Then, this implies that the least favorable distribution, which has at most $(k+1)(n+1)N$ point masses, must have a bounded support. Since the support is bounded, we can refine the bound on the number of mass points to $(n+1)(N-1) + k + 1$.

This concludes the proof.

B PROOF OF PROPOSITION 2

In this section, we present the detailed proof of Proposition 2. Choosing $\phi(x) = x^2$ generates the canonical example of a Bregman divergence, namely

$$\ell_\phi(X, f(Y)) = (X - f(Y))^2. \tag{45}$$

With part 5 of Theorem 1, the corresponding Bayesian risk is then of the form

$$R_\phi(P_X, P_{Y|X}) = \mathbb{E}[\ell_\phi(X, \mathbb{E}[X|Y])] = \mathbb{E}[(X - \mathbb{E}[X|Y])^2] =: g(\mathbf{x}), \tag{46}$$

which is nothing but the minimum mean square error. In what follows, we assume that g is differentiable.

To compute the gradient of g , the following well-known formula will be useful:

$$\mathbb{E}[(X - \mathbb{E}[X|Y])^2] = \mathbb{E}[X^2] - \mathbb{E}[\mathbb{E}[X|Y]^2]. \quad (47)$$

B.1 Partial Derivatives with Respect to Probabilities

In this subsection we focus on finding the partial derivatives of (46) with respect to the probability masses p_i , $i = 1, \dots, N$. Towards this end, consider the first term on the right-hand side of (47) and observe that

$$\frac{\partial}{\partial p_i} \mathbb{E}[X^2] = \frac{\partial}{\partial p_i} \sum_{k=1}^N p_k x_k^2 = x_i^2, \quad i = 1, \dots, N. \quad (48)$$

In order to find the partial derivative of the second term (i.e., $\partial \mathbb{E}[\mathbb{E}^2[X|Y]]/\partial p_i$), note first that

$$\frac{\partial}{\partial p_i} P_Y(y_j) = \frac{\partial}{\partial p_i} \sum_{k=1}^N p_k P_{Y|X}(y_j|x_k) = P_{Y|X}(y_j|x_i), \quad i = 1, \dots, N. \quad (49)$$

Furthermore, note that for $i = 1, \dots, N$

$$\frac{\partial}{\partial p_i} \mathbb{E}[X|Y = y_j] = \frac{\partial}{\partial p_i} \frac{\mathbb{E}[X P_{Y|X}(y_j|X)]}{P_Y(y_j)} \quad (50)$$

$$= \frac{P_Y(y_j) \frac{\partial}{\partial p_i} \mathbb{E}[X P_{Y|X}(y_j|X)] - \mathbb{E}[X P_{Y|X}(y_j|X)] \frac{\partial}{\partial p_i} P_Y(y_j)}{(P_Y(y_j))^2} \quad (51)$$

$$= \frac{P_Y(y_j) x_i P_{Y|X}(y_j|x_i) - \mathbb{E}[X P_{Y|X}(y_j|X)] P_{Y|X}(y_j|x_i)}{(P_Y(y_j))^2} \quad (52)$$

$$= \frac{x_i P_{Y|X}(y_j|x_i)}{P_Y(y_j)} - \mathbb{E}[X|Y = y_j] \frac{P_{Y|X}(y_j|x_i)}{P_Y(y_j)} \quad (53)$$

$$= (x_i - \mathbb{E}[X|Y = y_j]) \frac{P_{Y|X}(y_j|x_i)}{P_Y(y_j)}, \quad (54)$$

where in (51) we have used the derivative in (49) and

$$\frac{\partial}{\partial p_i} \mathbb{E}[X P_{Y|X}(y_j|X)] = \frac{\partial}{\partial p_i} \sum_{k=1}^N p_k x_k P_{Y|X}(y_j|x_k) = x_i P_{Y|X}(y_j|x_i). \quad (55)$$

Now,

$$\frac{\partial}{\partial p_i} \mathbb{E}[\mathbb{E}[X|Y]^2] = \frac{\partial}{\partial p_i} \sum_{j=1}^N P_Y(y_j) \mathbb{E}[X|Y = y_j]^2 \quad (56)$$

$$= \sum_{j=1}^N \mathbb{E}[X|Y = y_j]^2 \frac{\partial}{\partial p_i} P_Y(y_j) + \sum_{j=1}^N P_Y(y_j) \frac{\partial}{\partial p_i} \mathbb{E}[X|Y = y_j]^2 \quad (57)$$

$$= \sum_{j=1}^N P_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j]^2 + 2 \sum_{j=1}^N P_Y(y_j) \mathbb{E}[X|Y = y_j] (x_i - \mathbb{E}[X|Y = y_j]) \frac{P_{Y|X}(y_j|x_i)}{P_Y(y_j)} \quad (58)$$

$$= \mathbb{E}[\mathbb{E}[X|Y]^2 | X = x_i] + 2\mathbb{E}[\mathbb{E}[X|Y](X - \mathbb{E}[X|Y]) | X = x_i] \quad (59)$$

$$= 2x_i \mathbb{E}[\mathbb{E}[X|Y] | X = x_i] - \mathbb{E}[\mathbb{E}[X|Y]^2 | X = x_i], \quad (60)$$

where in (58) we have used expressions (49) and (54).

Finally, combining (48) and (60) results in

$$\frac{\partial}{\partial p_i} \mathbb{E}[(X - \mathbb{E}[X|Y])^2] = x_i^2 - 2x_i \mathbb{E}[\mathbb{E}[X|Y] | X = x_i] + \mathbb{E}[\mathbb{E}[X|Y]^2 | X = x_i] \quad (61)$$

$$= \mathbb{E}[(x_i - \mathbb{E}[X|Y])^2 | X = x_i], \quad (62)$$

$i = 1, \dots, N$.

B.2 Partial Derivatives with Respect to Locations

Now, we focus on finding the partial derivatives of (46) with respect to the locations x_i , $i = 1, \dots, N$, of the probability masses. Therefore, consider again the first term on the right-hand side of (47):

$$\frac{\partial}{\partial x_i} \mathbb{E}[X^2] = \frac{\partial}{\partial x_i} \sum_{k=1}^N p_k x_k^2 = 2x_i p_i, \quad i = 1, \dots, N. \quad (63)$$

To obtain the partial derivatives of the second term (47), first observe that

$$\frac{\partial}{\partial x_i} P_Y(y_j) = \frac{\partial}{\partial x_i} \sum_{k=1}^N p_k P_{Y|X}(y_j|x_k) = p_i \frac{\partial}{\partial x_i} P_{Y|X}(y_j|x_i) = p_i P'_{Y|X}(y_j|x_i). \quad (64)$$

Furthermore, we need the partial derivatives of the conditional expectation with respect to the locations, which is given by

$$\frac{\partial}{\partial x_i} \mathbb{E}[X|Y = y_j] = \frac{P_Y(y_j) \frac{\partial}{\partial x_i} \mathbb{E}[X P_{Y|X}(y_j|X)] - \mathbb{E}[X P_{Y|X}(y_j|X)] \frac{\partial}{\partial x_i} P_Y(y_j)}{(P_Y(y_j))^2} \quad (65)$$

$$= p_i \frac{x_i P'_{Y|X}(y_j|x_i) + P_{Y|X}(y_j|x_i)}{P_Y(y_j)} - \mathbb{E}[X|Y = y_j] \frac{p_i P'_{Y|X}(y_j|x_i)}{P_Y(y_j)} \quad (66)$$

$$= p_i \frac{(x_i - \mathbb{E}[X|Y = y_j]) P'_{Y|X}(y_j|x_i) + P_{Y|X}(y_j|x_i)}{P_Y(y_j)}, \quad (67)$$

where in (66) we have used the partial derivative (64) together with

$$\frac{\partial}{\partial p_i} \mathbb{E}[X P_{Y|X}(y_j|X)] = \frac{\partial}{\partial p_i} \sum_{k=1}^N p_k x_k P_{Y|X}(y_j|x_k) = p_i x_i P'_{Y|X}(y_j|x_i) + p_i P_{Y|X}(y_j|x_i). \quad (68)$$

Now, following along similar lines as in the previous subsection we obtain

$$\frac{\partial}{\partial x_i} \mathbb{E}[\mathbb{E}[X|Y]^2] = \frac{\partial}{\partial x_i} \sum_{j=1}^N P_Y(y_j) \mathbb{E}[X|Y = y_j]^2 \quad (69)$$

$$= \sum_{j=1}^N \mathbb{E}[X|Y = y_j]^2 \frac{\partial}{\partial x_i} P_Y(y_j) + 2 \sum_{j=1}^N P_Y(y_j) \mathbb{E}[X|Y = y_j] \frac{\partial}{\partial x_i} \mathbb{E}[X|Y = y_j] \quad (70)$$

$$= \sum_{j=1}^N p_i P'_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j]^2 + 2p_i \sum_{j=1}^N \mathbb{E}[X|Y = y_j] (x_i - \mathbb{E}[X|Y = y_j]) P'_{Y|X}(y_j|x_i) + P_{Y|X}(y_j|x_i) \quad (71)$$

$$= 2p_i \sum_{j=1}^N P_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j] + 2p_i x_i \sum_{j=1}^N P'_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j] \quad (72)$$

$$- \sum_{j=1}^N p_i P'_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j]^2. \quad (73)$$

Finally, combining (63) and (73) results in

$$\frac{\partial}{\partial x_i} \mathbb{E}[(X - \mathbb{E}[X|Y])^2] = 2x_i p_i - 2p_i \sum_{j=1}^N P_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j] - 2p_i x_i \sum_{j=1}^N P'_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j] \quad (74)$$

$$+ \sum_{j=1}^N p_i P'_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j]^2 \quad (75)$$

$$= 2p_i \left(x_i - \sum_{j=1}^N P_{Y|X}(y_j|x_i) \mathbb{E}[X|Y = y_j] \right) + p_i \sum_{j=1}^N P'_{Y|X}(y_j|x_i) (\mathbb{E}[X|Y = y_j]^2 \quad (76)$$

$$- 2x_i \mathbb{E}[X|Y = y_j]) \quad (77)$$

$$= 2p_i \left(x_i - \mathbb{E}[\mathbb{E}[X|Y] | X = x_i] \right) + p_i \sum_{j=1}^N P'_{Y|X}(y_j|x_i) (\mathbb{E}[X|Y = y_j]^2 - 2x_i \mathbb{E}[X|Y = y_j]), \quad (78)$$

which concludes the proof.