

Johannes Rauh; František Matúš
Maximizing the Bregman divergence from a Bregman family

Kybernetika, Vol. 56 (2020), No. 5, 875–885

Persistent URL: <http://dml.cz/dmlcz/148488>

Terms of use:

© Institute of Information Theory and Automation AS CR, 2020

Institute of Mathematics of the Czech Academy of Sciences provides access to digitized documents strictly for personal use. Each copy of any part of this document must contain these *Terms of use*.



This document has been digitized, optimized for electronic delivery and stamped with digital signature within the project *DML-CZ: The Czech Digital Mathematics Library* <http://dml.cz>

MAXIMIZING THE BREGMAN DIVERGENCE FROM A BREGMAN FAMILY

JOHANNES RAUH AND FRANTIŠEK MATUŠ

The problem to maximize the information divergence from an exponential family is generalized to the setting of Bregman divergences and suitably defined Bregman families.

Keywords: Bregman divergence, relative entropy, exponential family, optimization

Classification: 94A17, 62B05, 62E15, 62E17, 52A41

1. INTRODUCTION

Let Z be a finite set. Denote by $\mathcal{P}(Z)$ the set of probability measures (pm's) with support contained in Z . Let $\mathcal{E} \subseteq \mathcal{P}(Z)$ be an exponential family supported on Z . For $P, Q \in \mathcal{P}(Z)$ denote by $D(P\|Q)$ the *information divergence* (also known as *Kullback–Leibler divergence*), and let $D(P\|\mathcal{E}) \triangleq \inf_{Q \in \mathcal{E}} D(P\|Q)$. In 2002, Nihat Ay formulated the following optimization problem [1]:

Problem 1.1. Maximize $D(P\|\mathcal{E})$ over all probability distributions P on Z .

The original motivation came from theoretical studies of the infomax principle. Insight into this problem can also be used to bound approximation errors of machine learning models or other statistical models [13, 14].

Since 2002, progress has been made in different directions. The problem was attacked for particular classes of exponential families, with a particular focus on hierarchical models [2, 7, 9, 10]. A full characterization of the first order optimality conditions was given in [8].

In 2010, the first author found a surprising connection to another optimization problem [15]: Let A be the *design matrix* (or *sufficient statistics matrix*) of \mathcal{E} , where the columns of A are indexed by Z . Any $u \in \ker A$ can be written uniquely as a difference $u = u^+ - u^-$ of non-negative vectors u^+, u^- of disjoint support. For $u \in \ker A \setminus \{0\}$ with $\sum_{x \in Z} u^+(x) = \sum_{x \in Z} u^-(x) = 1$ let

$$\overline{D}(u) = H(u^-) - H(u^+) = \sum_{x \in Z} u(x) \log |u(x)|,$$

where H denotes the (Shannon) entropy. The second optimization problem is:

Problem 1.2. Maximize $\bar{D}(u)$ over the set all $u \in \ker A$ that satisfy $\sum_{x \in Z} u^+(x) = \sum_{x \in Z} u^-(x) = 1$.

The optimization problem 1.2 is easier than the optimization problem 1.1, since the function to be optimized in 1.1 is itself defined by an optimization problem.

Both authors showed in [12] that the map $u \mapsto u^+$ induces a one-to-one correspondence between the points that satisfy the respective critical equations of 1.2 and 1.1, and that this correspondence restricts to bijections of the sets of local optimizers and global optimizers, respectively.

The authors found this connection quite surprising. To better understand this result, the second author suggested to try to generalize the result to the setting of Bregman divergences and Bregman families. The present paper summarizes the results of this investigation.

The first step is the definition of a function \bar{B} that serves as an analogue of \bar{D} in the general case. Once this definition is in place, the equivalence of the global maximizers is rather straightforward (Theorem 5.1). What makes the general Bregman case more difficult is that \bar{B} is only defined implicitly as a solution of an optimization problem. Hence, the criticality conditions of \bar{B} are currently unknown. If the optimization problem underlying \bar{B} always has a unique solution (Conjecture 4.3), then the bijection of the local maximizers also generalizes (Theorem 5.5).

Section 2 recalls definitions and basic properties of Bregman divergences and introduces Bregman families. Section 3 discusses the problem of maximizing the Bregman divergence from a Bregman family. Section 4 introduces the function \bar{B} that corresponds to the function \bar{D} . Section 5 contains the main results that relates the problems to maximize the Bregman divergence and \bar{B} , respectively. Section 6 compares the results to the results of [12] that concern the classical case of exponential families and the information divergence.

2. PRELIMINARIES: BREGMAN DIVERGENCES AND BREGMAN FAMILIES

This section summarizes the relevant results about Bregman divergences and Bregman families. The end of the section contains in Example 2.7 the special case of information divergence and exponential families. For more details and generalizations to the case where Z is not finite see [11].

It is wellknown that one can associate to each exponential family a Bregman divergence by expressing the information divergence within the exponential family in terms of the exponential family's natural parameters. However, this construction is not used in this paper. Instead, starting from a particular Bregman divergence, a family of distributions is defined, called a *Bregman family*. These Bregman families generalize exponential families.

Consider a finite set Z . For each $z \in Z$ let $\beta_z : (0, +\infty) \rightarrow \mathbb{R}$ be a convex differentiable function with $\lim_{x \rightarrow 0^+} \beta'_z(x) = -\infty$ and $\lim_{x \rightarrow +\infty} \beta'_z(x) = +\infty$, where $\beta'_z(x)$ denotes the derivative of $\beta_z(x)$ with respect to x . Then the convex conjugate (see [17])

$$\beta_z^*(t) = \sup_x \{tx - \beta_z(x)\}$$

is differentiable and ranges between $-\lim_{x \rightarrow 0^+} \beta_z(x)$ and $+\infty$. The derivative $e_z(x) \triangleq \beta_z^*(x)$ is continuous and strictly increases from 0 to $+\infty$. Therefore, the inverse function $l_z(y) \triangleq e_z^{-1}(y)$ exists for $0 < y < +\infty$, is continuous and strictly increases from $-\infty$ to $+\infty$. The inverse function satisfies $l_z(y) = \beta'_z(y)$.

The following lemma is a standard result in convex analysis (see [17] or Lemma 2.2 in [11]):

Lemma 2.1. $\beta_z(e_z(r)) = re_z(r) - \beta^*(r)$ for all $r < \beta'(+\infty)$.

Consider a function $f : Z \rightarrow \mathbb{R}^d$. For $\vartheta \in \mathbb{R}^d$ define a pm $P_\vartheta : z \mapsto e_z(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta))$, where $\Lambda(\vartheta)$ is the unique solution of $\sum_{z \in Z} e_z(\langle \vartheta, f(z) \rangle - r) = 1$ in r . The subset

$$\mathcal{E} = \mathcal{E}_f \triangleq \{P_\vartheta : \vartheta \in \mathbb{R}^d\}$$

of $\mathcal{P}(Z)$ will be called a *Bregman family* in the following.¹ The matrix A with columns $f(z)$ for $z \in Z$ (after fixing an ordering of Z) is called the *design matrix* of \mathcal{E} .

The set $\text{cs}(\mathcal{E}) \triangleq \text{conv}\{f(z) : z \in Z\}$ is called the *convex support* of \mathcal{E} . The convex support is a (convex) polytope. A set $S \subseteq Z$ is called *facial* for \mathcal{E} if and only if $\text{conv}\{f(z) : z \in S\}$ is a face of $\text{cs}(\mathcal{E})$.

The *Bregman divergence* of $u, v : Z \rightarrow [0, +\infty)$ is

$$B(u, v) = \sum_{z \in Z} [\beta_z(u(z)) - \beta_z(v(z)) - \beta'_z(v(z))[u(z) - v(z)]]$$

The Bregman divergence of $P \in \mathcal{P}(Z)$ from a Bregman family \mathcal{E} is

$$B(P, \mathcal{E}) \triangleq \inf_{Q \in \mathcal{E}} B(P, Q)$$

When the minimizer in the definition of $B(P, \mathcal{E})$ does not exist, one can find a minimizer in the closure $\bar{\mathcal{E}}$ of \mathcal{E} , where the closure can be taken with respect to the canonical topology on the finite dimensional convex polytope $\mathcal{P}(Z)$. Just as in the classical case of an exponential family, one can prove the following statements:

Proposition 2.2. Let \mathcal{E} be a Bregman family.

1. For any $P \in \mathcal{P}(Z)$ there exists a unique pm $\Pi_{\mathcal{E}, P} \in \bar{\mathcal{E}}$ with

$$B(P, \Pi_{\mathcal{E}, P}) = B(P, \mathcal{E})$$

2. Let $P \in \mathcal{P}(Z)$ and $Q \in \bar{\mathcal{E}}$. If $\mathbb{E}_P[f] = \mathbb{E}_Q[f]$, then $Q = \Pi_{\mathcal{E}, P}$.
3. Let $P \in \mathcal{P}(Z)$. The unique global minimum of $H(Q) \triangleq \sum_{z \in Z} \beta_z(Q(z))$ for pm's $Q \in \mathcal{P}(Z)$ with $\mathbb{E}_P[f] = \mathbb{E}_Q[f]$ is given by $Q = \Pi_{\mathcal{E}, P}$.
4. The support $s(\Pi_{\mathcal{E}, P})$ is the smallest facial set containing $s(P)$.

¹The second author had originally given the name *generalized exponential family* to \mathcal{E} , which is also used by other authors. However, since that name is not very specific and since there are many different ways in which exponential families can be generalized, this paper now uses the name *Bregman family*.

The pm $\Pi_{\mathcal{E},P}$ is called the *generalized reverse Bregman projection* (rB -projection) of P to \mathcal{E} . Here, “generalized” may be dropped whenever $\Pi_{\mathcal{E},P} \in \mathcal{E}$. If the Bregman family \mathcal{E} is clear from the context, $\Pi_{\mathcal{E},P}$ is abbreviated by Π_P .

Proposition 2.3. Let \mathcal{E} be a Bregman family.

1. The map $\mu : P \in \mathcal{P}(Z) \mapsto \mathbb{E}_P[f]$ surjects onto $\text{cs}(\mathcal{E})$. It restricts to a homeomorphism $\bar{\mathcal{E}} \cong \text{cs}(\mathcal{E})$.
2. $\bar{\mathcal{E}} = \bigcup_F \mathcal{E}_F$, where F runs over all sets $F \subseteq Z$ that are facial with respect to \mathcal{E} and where \mathcal{E}_F is the Bregman family defined on F using $f|_F$.

For exponential families, the statements in Propositions 2.2 and 2.3 are well-known and go back at least to [3]. The statements continue to hold for exponential families when Z is replaced by a more general measure spaces Z , as studied in [4, 5]. The extended arXiv version of [18] contains a direct proof of the discrete case, which relies on algebraic insights from [6].

For a distribution of the form $Q : z \mapsto e_z(r_z)$, with $r_z \in \mathbb{R}$, by Lemma 2.1,

$$\begin{aligned} B(P, Q) &= \sum_{z \in Z} [\beta_z(P(z)) - \beta_z(e_z(r_z)) - r_z[P(z) - Q(z)]] \\ &= \sum_{z \in Z} [\beta_z(P(z)) - r_z e_z(r_z) + \beta^*(r_z) - r_z[P(z) - Q(z)]] \\ &= \sum_{z \in Z} [\beta_z(P(z)) + \beta^*(r_z) - r_z P(z)]. \end{aligned}$$

When $Q \in \mathcal{E}$, then r_z is of the the form $\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)$. Thus,

$$\begin{aligned} B(P, \mathcal{E}) &= \sum_z \beta_z(P(z)) - \sup_{\vartheta} \left[\langle \vartheta, \sum_z f(z)P(z) \rangle - \Lambda(\vartheta) - \sum_z \beta_z^*(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) \right] \\ &= \sum_z \beta_z(P(z)) - \sup_{\vartheta} \left[\langle \vartheta, \mu(P) \rangle - \Upsilon(\vartheta) \right], \quad (1) \end{aligned}$$

where $\Upsilon(\vartheta) = \Lambda(\vartheta) + \sum_z \beta_z^*(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta))$.

Theorem 2.4. Υ is convex. Its partial derivatives are

$$\frac{\partial}{\partial \vartheta_i} \Upsilon(\vartheta) = \mathbb{E}_{\vartheta}[f_i] = \mu(P_{\vartheta})_i,$$

where \mathbb{E}_{ϑ} denotes the expected value taken with respect to P_{ϑ} . The map $\nabla \Upsilon : \mathbb{R}^d \rightarrow \text{cs}(\mathcal{E})$ is surjective. The Hessian of Υ is positive definite.

Proof.

$$\begin{aligned} \frac{\partial}{\partial \vartheta_i} \Upsilon(\vartheta) &= \partial_i \Lambda(\vartheta) + \sum_z \beta_z^{*'}(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) [f_i(z) - \partial_i \Lambda(\vartheta)] \\ &= \sum_z \beta_z^{*'}(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) f_i(z), \\ \frac{\partial^2}{\partial \vartheta_i \partial \vartheta_j} \Upsilon(\vartheta) &= \sum_z \beta_z^{*''}(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) f_i(z) [f_j(z) - \partial_j \Lambda(\vartheta)] \\ &= \sum_z \beta_z^{*''}(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) [f_i(z) - \partial_i \Lambda(\vartheta)] [f_j(z) - \partial_j \Lambda(\vartheta)] \succeq 0, \end{aligned}$$

where the last equality follows from taking the derivative of the defining equation $\sum_z \beta_z^{*'}(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) = 1$ of $\Lambda(\vartheta)$:

$$0 = \frac{\partial}{\partial \vartheta_j} \sum_z \beta_z^{*'}(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) = \sum_z \beta_z^{*''}(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) [f_j(z) - \frac{\partial}{\partial \vartheta_j} \Lambda(\vartheta)].$$

This shows convexity.

It is clear that $\mathbb{E}_\vartheta[f] = \mu(P_\vartheta)$ belongs to $\text{conv}\{f(z) : z \in Z\}$. Surjectivity follows from Proposition 2.3. □

It follows from the properties of convex conjugation:

Corollary 2.5. The maps $\vartheta \mapsto \nabla \Upsilon(\vartheta)$ and $\mu \mapsto \nabla \Upsilon^*(\mu)$ are mutual inverses in the relative interiors of their respective domains. If $\Pi_P \in \text{ri}(\mathcal{E})$, then $\Pi_P = P_\vartheta$ for $\vartheta = \nabla \Upsilon^*(\mu(P))$.

Let $H(P) = \sum_{z \in Z} \beta_z(P(z))$ as in Proposition 2.2. Then (1) rewrites to

$$B(P, \mathcal{E}) = H(P) - \Upsilon^*(\mu(P)),$$

where Υ^* denotes the convex conjugate of Υ . From this equality follows the next result, which can also be seen as a kind of Pythagorean identity:

Corollary 2.6. $B(P, \mathcal{E}) = H(P) - H(\Pi_P)$ for all $P \in \mathcal{P}(Z)$.

Proof. $B(P, \mathcal{E}) = B(P, \mathcal{E}) - B(\Pi_P, \mathcal{E}) = H(P) - H(\Pi_P)$, since $\mu(P) = \mu(\Pi_P)$. □

Example 2.7. Let $\beta_z(x) = x \ln(x/\nu(z)) + x$ for all $z \in Z$. Then $\beta_z^*(x) = \nu(z) \exp(x)$ and $l_z(x) = \beta_z'(x) = \ln(x/\nu(z))$, and so $e_z(x) = \nu(z) \exp(x)$. In this case, \mathcal{E} is an exponential family with reference measure ν , and B equals the information divergence. Since $\beta_z^* = e_z(x)$, it follows that $\sum_z \beta_z^*(\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)) = 1$. Therefore, $\Upsilon(\vartheta) = 1 + \Lambda(\vartheta)$. In the classical case, Λ is called the *partition function*, and convexity of Λ is well-known and widely used. In the general case, Λ itself need not be convex.

3. MAXIMIZING THE BREGMAN DIVERGENCE FROM A BREGMAN FAMILY

Let \mathcal{E} be a Bregman family. The following problem generalizes Problem 1.1:

Problem 3.1. Maximize $B(P, \mathcal{E})$ over $P \in \mathcal{P}(Z)$.

Theorem 3.2. If $P \in \mathcal{P}(Z)$ is a local maximizer of $B(\cdot, \mathcal{E})$, then the map $z \mapsto l(P(z)) - l(\Pi_P(z))$ is constant for $z \in s(P)$.

Proof. If $\mu(P) = \sum_z f(z)P(z)$ does not lie in the relative interior of $\text{cs}(\mathcal{E})$, by Proposition 2.3, one may replace \mathcal{E} by \mathcal{E}_F for some suitable $F \subsetneq Z$. Thus, without loss of generality, assume that $\mu(P)$ lies in the relative interior of $\text{cs}(\mathcal{E})$.

Let $w \in \mathbb{R}^Z$ with $\sum_z w(z) = 0$ and $s(w) \subseteq s(P)$. For $\epsilon > 0$ small,

$$B(P + \epsilon w, \mathcal{E}) \approx H(P) + \epsilon \sum_z \beta'_z(P(z))w(z) - \Upsilon^*(\mu(P)) - \epsilon \left\langle \nabla \Upsilon^*(\mu(P)), \sum_z f(z)w(z) \right\rangle$$

to first order in ϵ . Let $\vartheta = \nabla \Upsilon^*(\mu(P))$. Then $\Pi_P = P_\vartheta$ by Corollary 2.5, and

$$\left\langle \vartheta, \sum_z f(z)w(z) \right\rangle = \sum_z [\langle \vartheta, f(z) \rangle - \Lambda(\vartheta)]w(z) = \sum_z \beta'_z(\Pi_P(z))w(z),$$

since β' and $\beta^{*'}$ are mutual inverses to each other. In total,

$$B(P + \epsilon w, \mathcal{E}) \approx H(P) - \Upsilon^*(\mu(P)) + \epsilon \sum_z [\beta'_z(P(z)) - \beta'_z(\Pi_P(z))]w(z),$$

whence $\sum_z [\beta'_z(P(z)) - \beta'_z(\Pi_P(z))]w(z) = 0$ if P is a critical point. This equality holds for all $w \in \mathbb{R}^Z$ with $\sum_z w(z) = 0$ and $s(w) \subseteq s(P)$. Therefore, $\beta'_z(P(z)) - \beta'_z(\Pi_P(z))$ is constant for $z \in s(P)$. \square

Corollary 3.3. Let $P \in \mathcal{P}(Z)$ be a local maximizer of $B(\cdot, \mathcal{E})$, and let $u = P - \Pi_P$. Then $s(u^+) = s(P)$.

If $\beta_x = \beta_y$ for $x, y \in s(P)$, then $u^+(x) \geq u^+(y)$ if and only if $P(x) \geq P(y)$.

Proof. By Theorem 3.2, there exists a constant c such that $l(P(z)) - l(\Pi_P(z)) = c$ for $z \in s(P)$. The number c equals the unique solution of the equation

$$\sum_{z \in s(P)} e_z(l_z(\Pi_P(z)) + c) = \sum_{z \in s(P)} P(z) = 1.$$

Since all functions e_z are increasing, $c > 0$. Thus, if $z \in s(P)$, then $l_z(P(z)) > l_z(\Pi_P(z))$, and so $P(z) > \Pi_P(z)$. This implies $s(P) \subseteq s(u^+)$. On the other hand, if $z \notin s(P)$, then $\Pi_P(z) \geq P(z)$, and so $u(z) \leq 0$, which implies $s(P) \supseteq s(u^+)$. \square

As in the classical case, one shows [10]:

Proposition 3.4. Any $P \in \mathcal{P}(Z)$ that globally maximizes $B(\cdot, \mathcal{E})$ satisfies $|s(P)| \leq \dim(\mathcal{E}) + 1$.

4. THE FUNCTION \bar{B} AND THE ALTERNATIVE OPTIMIZATION PROBLEM

For each real vector-valued function $f : Z \rightarrow \mathbb{R}^d$ let

$$\mathcal{N} = \mathcal{N}(f) = \left\{ u \in \mathbb{R}^Z : \sum_{z \in Z} f(z)u(z) = 0 \text{ and } \sum_{z \in Z} u(z) = 0 \right\}.$$

If A is a design matrix, then $\mathcal{N} = \{u \in \ker A : \sum_{z \in Z} u(z) = 0\}$.

Let $u : Z \rightarrow \mathbb{R}$ be a real function satisfying $\sum_{x \in Z} u(x) = 0$. To each such u associate a function f_u such that $\mathcal{N}(f_u) = \mathbb{R}u$, and let $\mathcal{F}_u = \mathcal{E}_{f_u}$. Then \mathcal{F}_u has codimension one. By Proposition 2.2, the difference $P - \Pi_{\mathcal{F}_u, P}$ lies in $\mathbb{R}u$.

Lemma 4.1. Let $P \in \mathcal{P}(Z)$, and let $u = P - \Pi_P$. Then $\Pi_{\mathcal{F}_u, P} = \Pi_P$.

Proof. From $\mathcal{E} \subseteq \mathcal{F}_u$ follows $\Pi_P \in \mathcal{F}_u$. Together with $P - \Pi_P = u$, the statement follows from Proposition 2.2. □

$\mathcal{P}(Z)$ can be partitioned into $\mathcal{P}_u^+ \cup \bar{\mathcal{F}}_u \cup \mathcal{P}_u^-$, where

$$\mathcal{P}_u^+ = \{P \in \mathcal{P}(Z) : \langle P - \Pi_{\mathcal{F}_u, P}, u \rangle > 0\}, \quad \mathcal{P}_u^- = \{P \in \mathcal{P}(Z) : \langle P - \Pi_{\mathcal{F}_u, P}, u \rangle < 0\}.$$

The definition and Lemma 4.1 imply:

Lemma 4.2. $P \in \mathcal{P}_{P - \Pi_P}^+$ for any $P \in \mathcal{P}(Z) \setminus \bar{\mathcal{E}}$.

In the classical case, the maximizer of the information divergence from an arbitrary exponential family \mathcal{E} need not be unique [10]. However, when $\mathcal{E} = \mathcal{F}_u$ has codimension one, there are precisely two local maximizers u^+ and u^- , one on each side of \mathcal{E} [16, Section VI]. This motivates the following conjecture:

Conjecture 4.3. The map $P \in \mathcal{P}_u^+ \mapsto B(P, \mathcal{F}_u)$ has a unique local (and global) maximizer.

The proof of the conjecture in the classical case relies on applying properties of the logarithm to the criticality conditions in Theorem 3.2. It is not possible to apply this proof to the general case of the conjecture.

For any function $u : Z \rightarrow \mathbb{R}$ that satisfies $\sum_{z \in Z} u(z) = 0$ let

$$\bar{B}(u) \triangleq \max \{B(P, \mathcal{F}_u) : P \in \bar{\mathcal{P}}_u^+\},$$

where $\bar{\mathcal{P}}_u^+ = \mathcal{P}_u \cup \bar{\mathcal{F}}_u$ denotes the closure of \mathcal{P}_u^+ . The map \bar{B} is continuous and welldefined since $\bar{\mathcal{P}}_u^+ = \mathcal{P}_u^+ \cup \bar{\mathcal{F}}_u$ is compact. If $u \neq 0$, then this maximum lies in \mathcal{P}_u^+ , and $\bar{B}(u) > 0$. The function \bar{B} satisfies $\bar{B}(\lambda u) = \bar{B}(u)$ for all $\lambda > 0$.

Problem 4.4. Maximize the function $u \in \mathcal{N} \setminus \{0\} \mapsto \bar{B}(u)$.

The intuition behind the definition of \bar{B} and Problem 4.4 is the following: instead of directly searching for a maximizer P of $B(\cdot, \mathcal{E})$, one may try to determine the vector $u = P - \Pi_P$, which can be seen as a direction within the probability simplex. Thus, the task is to find a direction in which it is possible to achieve large values of $B(\cdot, \mathcal{E})$. When analyzing the direction u , Lemma 4.1 says that one may just as well replace \mathcal{E} by \mathcal{F}_u .

5. EQUIVALENCE OF THE MAXIMIZERS

The following theorem specifies the relations between the problems 3.1 and 4.4. It corresponds to [15, Theorem 3].

- Theorem 5.1.**
1. $\max_{P \in \mathcal{P}(Z)} B(P, \mathcal{E}) = \max_{u \in \mathcal{N} \setminus \{0\}} \bar{B}(u)$.
 2. If P is a global maximizer of problem 3.1, then $P - \Pi_P$ is a global maximizer of problem 4.4, and $B(P, \mathcal{E}) = \bar{B}(P - \Pi_P)$.
 3. If u is a global maximizer of problem 4.4 and if $\bar{B}(u) = B(P, \mathcal{F}_u)$, then P is a global maximizer of problem 3.1, and $\bar{B}(u) = B(P, \mathcal{E})$.

The proof of Theorem 5.1 is based on the following auxilliary theorem, which corresponds to [12, Theorem 2].

Theorem 5.2. $\bar{B}(P - \Pi_P) \geq B(P, \mathcal{E})$ for any $P \in \mathcal{P}(Z) \setminus \bar{\mathcal{E}}$. If $u \in \mathcal{N} \setminus \{0\}$ and $P \in \mathcal{P}(Z)$ satisfy $\bar{B}(u) = B(P, \mathcal{F}_u)$, then $B(P, \mathcal{E}) \geq \bar{B}(u)$, with equality if and only if $P - \Pi_P = \lambda u$ for some $\lambda > 0$.

Proof. The first statement follows from Lemma 4.1, as $\bar{B}(P - \Pi_P) \geq B(P, \mathcal{F}_{P - \Pi_P}) = B(P, \mathcal{E})$. For the second statement observe that from $\mathcal{E} \subseteq \mathcal{F}_u$ follows $B(P, \mathcal{E}) \geq B(P, \mathcal{F}_u) = \bar{B}(u)$. □

Theorem 5.1 follows directly from Theorem 5.2.

In [12, Theorem 1] it was shown that the points that satisfy the respective critical equations (i.e. the equality conditions among the first order conditions) of the two problems 3.1 and 4.4 and the local maximizers of the two problems are also in one-to-one correspondence in the classical case. Discussing the criticality conditions is difficult, as no explicit formula for \bar{B} is known, and if Conjecture 4.3 is wrong, it is improbable that \bar{B} is differentiable. If the conjecture is true, one can at least prove that the local maximizers of the two problems are related, as Theorem 5.5 below will show.

Assume that Conjecture 4.3 is true, and let $\Phi(u) \triangleq \arg \max_{Q \in \mathcal{P}_u^+(Z)} B(Q, \mathcal{F}_u)$ for $u \in \mathcal{N} \setminus \{0\}$. By assumption, Φ is well-defined and continuous. The map $\Psi : \mathcal{P} \rightarrow \mathcal{N}, P \mapsto P - \Pi_P$ is also continuous. With these two maps, Theorem 5.2 can be reformulated as follows:

Corollary 5.3. If Conjecture 4.3 is true, then:

1. $\bar{B}(\Psi(P)) \geq B(P, \mathcal{E})$ for all $P \in \mathcal{P}(Z)$, with equality if and only if $P = \Phi(\Psi(P))$.
2. $B(\Phi(u)) \geq \bar{B}(u)$ for all $u \in \mathcal{N} \setminus \{0\}$, with equality if and only if $\Psi(\Phi(u)) = \lambda u$ for some $\lambda > 0$.

Lemma 5.4.

1. If $u \in \mathcal{N} \setminus \{0\}$ is a local maximizer of \bar{B} , then $\Pi_{\Phi(u)} = \Pi_{\mathcal{F}_u, \Phi(u)}$. Thus, if Conjecture 4.3 is true, then $\Psi(\Phi(u)) = \lambda u$ for some $\lambda > 0$.

2. If $P \in \mathcal{P}(Z)$ is a local maximizer of $B(\cdot, \mathcal{E})$, then P is a local maximizer of $B(\cdot, \mathcal{F}_{\Psi(P)})$. Thus, if Conjecture 4.3 is true, then $\Phi(\Psi(P)) = P$.

Proof. For the first statement, let $P = \Phi(u)$. Suppose that $\Pi_P \neq \Pi_{\mathcal{F}_u, P}$, and let Q be a pm in the convex hull of Π_P and $\Pi_{\mathcal{F}_u, P}$. Since H is strictly convex and by Proposition 2.2, $H(\Pi_P) < H(Q) < H(\Pi_{\mathcal{F}_u, P})$. Corollary 5.3 implies $\overline{B}(P - Q) \geq B(P, \mathcal{F}_{P-Q}) \geq H(P) - H(Q) > H(P) - H(\Pi_{\mathcal{F}_u, P}) = \overline{B}(u)$. This contradicts the assumption that u is a local maximizer. Hence, $\Pi_{\mathcal{F}_u, P} = \Pi_P$, and $u = \Phi(P)$.

For any $Q \in \mathcal{P}(Z)$, if $B(Q, \mathcal{E}) \leq B(P, \mathcal{E})$, then $B(Q, \mathcal{F}_{P-\Pi_P}) \leq B(Q, \mathcal{E}) \leq B(P, \mathcal{E}) = B(P, \mathcal{F}_{P-\Pi_P})$, where the last equality uses Lemma 4.1. This proves the second statement. □

Theorem 5.5. Assume that Conjecture 4.3 is true. If P is a local maximizer of $B(P, \mathcal{E})$, then $\Psi(P)$ is a local maximizer of \overline{B} . If u is a local maximizer of \overline{B} , then $\Phi(u)$ is a local maximizer of $B(P, \mathcal{E})$.

Proof. Let P be a local maximizer of $B(P, \mathcal{E})$. Let U be a neighbourhood of P in $\mathcal{P}(Z)$ such that $B(Q, \mathcal{E}) \leq B(P, \mathcal{E})$ for all $Q \in U$. Then $U' := \Phi^{-1}(U)$ is a neighbourhood of P by Lemma 5.4. If $v \in U'$, then Corollary 5.3 implies

$$\overline{B}(\Psi(P)) \geq B(P, \mathcal{E}) \geq B(\Phi(v), \mathcal{E}) = \overline{B}(v).$$

for all $v \in U'$. This proves the first statement.

Let $u \in \mathcal{N} \setminus \{0\}$ be a local maximizer of \overline{B} . Let U' be a neighbourhood of u in $\mathcal{N} \setminus \{0\}$ with $\overline{B}(u) \geq \overline{B}(v)$ for all $v \in U'$. Then $U := \Psi^{-1}(U')$ is a neighbourhood of P by Lemma 5.4. If $Q \in U$, then Corollary 5.3 implies

$$B(\Phi(u), \mathcal{E}) \geq \overline{B}(u) \geq \overline{B}(\Psi(Q)) \geq B(Q, \mathcal{E}).$$

This proves the second statement. □

6. COMPARISON TO THE CLASSICAL CASE

In the classical case $\beta_z(t) = t \ln(t/\nu(z))$, in which B becomes the information (or Kullback–Leibler) divergence and \mathcal{E} is an exponential family with reference measure ν , the function \overline{B} , which, in the general case, is defined by means of an optimization problem, has an explicit analytic expression:

$$\overline{B}(u) = \ln \left(1 + \exp \left(\sum_{z \in Z} \frac{u(z)}{\|u\|_1} \ln |u(z)| \right) \right) = \ln (1 + \exp(\overline{D}(u))).$$

Thus, while an optimization problem has to be solved to evaluate the function $B(\cdot, \mathcal{E})$ at some $P \in \mathcal{P}(Z)$, the function \overline{B} can be evaluated more easily.

In the general case this is not true anymore. However, the computational complexity of the optimization problem 4.4 is still different from the complexity of the problem 3.1. To evaluate $\overline{B}(u)$ at a single point $u \in \mathcal{N} \setminus \{0\}$, a problem of a similar kind as problem 3.1, but much smaller, has to be solved: the solution is a pm in $\mathcal{P}(s(u^+))$. Moreover, as \mathcal{F}_u has co-dimension one, rB -projections to \mathcal{F}_u can be computed by solving a one-dimensional optimization problem (namely, $\Pi_{\mathcal{F}_u, P}$ minimizes $H(Q)$ for $Q \in P + \mathbb{R}u$).

In total, whether it is easier to attack problem 4.4 or 3.1 may depend on the specific choice of the functions β_z and f . For the classical case, [16] and [15] present many ideas how to attack problem 1.2, many of which may generalize to problem 4.4, depending on the choice of the functions β_z .

Most importantly, the idea behind the definition of the function \bar{B} sheds light on the relation of the problems 1.1 and 1.2, which is rather opaque if one only looks at the definitions of the functions D and \bar{D} .

ACKNOWLEDGEMENT

This work was partially supported by the Grant Agency of the Czech Republic under Grant P202/10/0618 and the Research Academy Leipzig.

AUTHOR CONTRIBUTIONS

The first investigations were done by the second author in 2010, who also provided the correct notion of a Bregman family. In 2012, both authors worked together to find a good definition for \bar{B} and to prove the equivalence of the global maximizers (Theorem 5.1). The project was delayed by the first author trying to find a proof of Conjecture 4.3. At the occasion of this special issue, first author added further results and completed the manuscript.

(Received April 10, 2019)

REFERENCES

-
- [1] N. Ay: An information-geometric approach to a theory of pragmatic structuring. *Ann. Probab.* *30* (2002), 416–436. DOI:10.1214/aop/1020107773
 - [2] N. Ay and A. Knauf: Maximizing multi-information. *Kybernetika* *42* (2006), 517–538.
 - [3] O. Barndorff-Nielsen: *Information and Exponential Families in Statistical Theory*. Wiley, 1978.
 - [4] I. Csiszár and F. Matúš: Closures of exponential families. *Ann. Probab.* *33* (2005), 582–600. DOI:10.1214/009117904000000766
 - [5] I. Csiszár and F. Matúš: Generalized maximum likelihood estimates for exponential families. *Probab. Theory Related Fields* *141* (2008), 213–246. DOI:10.1007/s00440-007-0084-z
 - [6] D. Geiger, C. Meek, and B. Sturmfels: On the toric algebra of graphical models. *Ann. Statist.* *34* (2006), 1463–1492. DOI:10.1214/009053606000000263
 - [7] F. Matúš: Maximization of information divergences from binary i.i.d. sequences. In: *Proc. IPMU 2* (2004), 1303–1306.
 - [8] F. Matúš: Optimality conditions for maximizers of the information divergence from an exponential family. *Kybernetika* *43* (2007), 731–746.
 - [9] F. Matúš: Divergence from factorizable distributions and matroid representations by partitions. *IEEE Trans. Inform. Theory* *55* (2009), 5375–5381. DOI:10.1109/tit.2009.2032806
 - [10] F. Matúš and N. Ay: On maximization of the information divergence from an exponential family. In: *Proc. WUPES 2003, University of Economics, Prague 2003*, pp. 199–204.
 - [11] F. Matúš and I. Csiszár: Generalized minimizers of convex integral functionals, Bregman distance, Pythagorean identities. *Kybernetika* *48* (2012), 637–689.

- [12] F. Matúš and J. Rauh: Maximization of the information divergence from an exponential family and criticality. In: Proc. IEEE International Symposium on Information Theory (ISIT2011), 2011. DOI:10.1109/isit.2011.6034269
- [13] G. Montúfar, J. Rauh, and N. Ay: Expressive power and approximation errors of Restricted Boltzmann Machines. In: Proc. NIPS 2011.
- [14] G. Montúfar, J. Rauh, and N. Ay: Maximal information divergence from statistical models defined by neural networks. In: Proc. GSI, 2013, pp. 759–766. DOI:10.1007/978-3-642-40020-9_85
- [15] J. Rauh: Finding the maximizers of the information divergence from an exponential family. IEEE Trans. Inform. Theory *57* (2011), 3236–3247. DOI:10.1109/tit.2011.2136230
- [16] J. Rauh: Finding the Maximizers of the Information Divergence from an Exponential Family Ph.D. Dissertation, Universität Leipzig, 2011.
- [17] R. T. Rockafellar: Convex Analysis. Princeton University Press, 1970. DOI:10.1017/s0013091500010142
- [18] N. Wang, J. Rauh, and H. Massam: Approximating faces of marginal polytopes in discrete hierarchical models. Ann. Statist. *47* (2019), 1203–1233. DOI:10.1214/18-aos1710. Extended preprint version at [arXiv:1603.04843](https://arxiv.org/abs/1603.04843).

*Johannes Rauh, Max Planck Institute for Mathematics in the Sciences, Inselstraße 22, 04103 Leipzig, Germany.
e-mail: jrauh@mis.mpg.de*