

Shun-ichi AMARI

Information geometry in optimization, machine learning and statistical inference

© Higher Education Press and Springer-Verlag Berlin Heidelberg 2010

Abstract The present article gives an introduction to information geometry and surveys its applications in the area of machine learning, optimization and statistical inference. Information geometry is explained intuitively by using divergence functions introduced in a manifold of probability distributions and other general manifolds. They give a Riemannian structure together with a pair of dual flatness criteria. Many manifolds are dually flat. When a manifold is dually flat, a generalized Pythagorean theorem and related projection theorem are introduced. They provide useful means for various approximation and optimization problems. We apply them to alternative minimization problems, Ying-Yang machines and belief propagation algorithm in machine learning.

Keywords information geometry, machine learning, optimization, statistical inference, divergence, graphical model, Ying-Yang machine

1 Introduction

Information geometry [1] deals with a manifold of probability distributions from the geometrical point of view. It studies the invariant structure by using the Riemannian geometry equipped with a dual pair of affine connections. Since probability distributions are used in many problems in optimization, machine learning, vision, statistical inference, neural networks and others, information geometry provides a useful and strong tool to many areas of information sciences and engineering.

Many researchers in these fields, however, are not familiar with modern differential geometry. The present

article intends to give an understandable introduction to information geometry without modern differential geometry. Since underlying manifolds in most applications are dually flat, the dually flat structure plays a fundamental role. We explain the fundamental dual structure and related dual geodesics without using the concept of affine connections and covariant derivatives.

We begin with a divergence function between two points in a manifold. When it satisfies an invariance criterion of information monotonicity, it gives a family of f -divergences [2]. When a divergence is derived from a convex function in the form of Bregman divergence [3], this gives another type of divergence, where the Kullback-Leibler divergence belongs to both of them. We derive a geometrical structure from a divergence function [4]. The Fisher information Riemannian structure is derived from an invariant divergence (f -divergence) (see Refs. [1,5]), while the dually flat structure is derived from the Bregman divergence (convex function).

The manifold of all discrete probability distributions is dually flat, where the Kullback-Leibler divergence plays a key role. We give the generalized Pythagorean theorem and projection theorem in a dually flat manifold, which plays a fundamental role in applications. Such a structure is not limited to a manifold of probability distributions, but can be extended to the manifolds of positive arrays, matrices and visual signals, and will be used in neural networks and optimization problems.

After introducing basic properties, we show three areas of applications. One is application to the alternative minimization procedures such as the expectation-maximization (EM) algorithm in statistics [6–8]. The second is an application to the Ying-Yang machine introduced and extensively studied by Xu [9–14]. The third one is application to belief propagation algorithm of stochastic reasoning in machine learning or artificial intelligence [15–17]. There are many other applications in analysis of spiking patterns of the brain, neural networks, boosting algorithm of machine learning, as well as wide range of statistical inference, which we do not

Received January 15, 2010; accepted February 5, 2010

Shun-ichi AMARI (✉)

RIKEN Brain Science Institute, Saitama 351-0198, Japan

E-mail: amari@brain.riken.jp

mention here.

2 Divergence function and information geometry

2.1 Manifold of probability distributions and positive arrays

We introduce divergence functions in various spaces or manifolds. To begin with, we show typical examples of manifolds of probability distributions. A one-dimensional Gaussian distribution with mean μ and variance σ^2 is represented by its probability density function

$$p(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left\{ -\frac{(x - \mu)^2}{2\sigma^2} \right\}. \quad (1)$$

It is parameterized by a two-dimensional parameter $\xi = (\mu, \sigma)$. Hence, when we treat all such Gaussian distributions, not a particular one, we need to consider the set S_G of all the Gaussian distributions. It forms a two-dimensional manifold

$$S_G = \{p(x; \xi)\}, \quad (2)$$

where $\xi = (\mu, \sigma)$ is a coordinate system of S_G . This is not the only coordinate system. It is possible to use other parameterizations or coordinate systems when we study S_G .

We show another example. Let x be a discrete random variable taking values on a finite set $X = \{0, 1, \dots, n\}$. Then, a probability distribution is specified by a vector $\mathbf{p} = (p_0, p_1, \dots, p_n)$, where

$$p_i = \text{Prob} \{x = i\}. \quad (3)$$

We may write

$$p(x; \mathbf{p}) = \sum p_i \delta_i(x), \quad (4)$$

where

$$\delta_i(x) = \begin{cases} 1, & x = i, \\ 0, & x \neq i. \end{cases} \quad (5)$$

Since \mathbf{p} is a probability vector, we have

$$\sum p_i = 1, \quad (6)$$

and we assume

$$p_i > 0. \quad (7)$$

The set of all the probability distributions is denoted by

$$S_n = \{\mathbf{p}\}, \quad (8)$$

which is an n -dimensional simplex because of (6) and (7). When $n = 2$, S_n is a triangle (Fig. 1). S_n is an n -dimensional manifold, and $\xi = (p_1, \dots, p_n)$ is a coordinate system. There are many other coordinate systems. For example,

$$\theta_i = \log \frac{p_i}{p_0}, \quad i = 1, \dots, n, \quad (9)$$

is an important coordinate system of S_n , as we will see later.




Fig. 1 Manifold S_2 of discrete probability distributions

The third example deals with positive measures, not probability measures. When we disregard the constraint $\sum p_i = 1$ of (6) in S_n , keeping $p_i > 0$, \mathbf{p} is regarded as an $(n + 1)$ -dimensional positive arrays, or a positive measure where $x = i$ has measure p_i . We denote the set of positive measures or arrays by

$$M_{n+1} = \{\mathbf{z}, z_i > 0; i = 0, 1, \dots, n\}. \quad (10)$$

This is an $(n + 1)$ -dimensional manifold with a coordinate system \mathbf{z} . S_n is its submanifold derived by a linear constraint $\sum z_i = 1$.

In general, we can regard any regular statistical model

$$S = \{p(x, \xi)\} \quad (11)$$

parameterized by ξ as a manifold with a (local) coordinate system ξ . It is a space M of positive measures, when the constraint $\int p(x, \xi) dx = 1$ is discarded. We may treat any other types of manifolds and introduce dual structures in them. For example, we will consider a manifold consisting of positive-definite matrices.

2.2 Divergence function and geometry

We consider a manifold S having a local coordinate system $\mathbf{z} = (z_i)$. A function $D[\mathbf{z} : \mathbf{w}]$ between two points \mathbf{z} and \mathbf{w} of S is called a divergence function when it satisfies the following two properties:

- 1) $D[\mathbf{z} : \mathbf{w}] \geq 0$, with equality when and only when $\mathbf{z} = \mathbf{w}$.
- 2) When the difference between \mathbf{w} and \mathbf{z} is infinitesimally small, we may write $\mathbf{w} = \mathbf{z} + d\mathbf{z}$ and Taylor expansion gives

$$D[\mathbf{z} : \mathbf{z} + d\mathbf{z}] = \sum g_{ij}(\mathbf{z}) dz_i dz_j, \quad (12)$$

where

$$g_{ij}(\mathbf{z}) = \frac{\partial^2}{\partial z_i \partial z_j} D[\mathbf{z} : \mathbf{w}]|_{\mathbf{w}=\mathbf{z}} \quad (13)$$

is a positive-definite matrix.

A divergence does not need to be symmetric, and

$$D[\mathbf{z} : \mathbf{w}] \neq D[\mathbf{w} : \mathbf{z}] \tag{14}$$

in general, nor does it satisfy the triangular inequality. Hence, it is not a distance. It rather has a dimension of square of distance as is seen from (12). So (12) is considered to define the square of the local distance ds between two nearby points $Z = (\mathbf{z})$ and $Z + dZ = (\mathbf{z} + d\mathbf{z})$,

$$ds^2 = \sum g_{ij}(\mathbf{z}) dz_i dz_j. \tag{15}$$

More precisely, $d\mathbf{z}$ is regarded as a small line element connecting two points Z and $Z + dZ$. This is a tangent vector at point Z (Fig. 2). When a manifold has a positive-definite matrix $g_{ij}(\mathbf{z})$ at each point \mathbf{z} , it is called a Riemannian manifold, and (g_{ij}) is a Riemannian metric tensor.




Fig. 2 Manifold, tangent space and tangent vector

An affine connection defines a correspondence between two nearby tangent spaces. By using it, a geodesic is defined: A geodesic is a curve of which the tangent directions do not change along the curve by this correspondence (Fig. 3). It is given mathematically by a covariant derivative, and technically by the Christoffel symbol, $\Gamma_{ijk}(\mathbf{z})$, which has three indices.




Fig. 3 Geodesic, keeping the same tangent direction

One may skip the following two paragraphs, since technical details are not used in the following. Eguchi [4] proposed the following two connections:

$$\Gamma_{ijk}(\mathbf{z}) = -\frac{\partial^3}{\partial z_i \partial z_j \partial w_k} D[\mathbf{z} : \mathbf{w}]|_{\mathbf{w}=\mathbf{z}}, \tag{16}$$

$$\Gamma_{ijk}^*(\mathbf{z}) = -\frac{\partial^3}{\partial w_i \partial w_j \partial z_k} D[\mathbf{z} : \mathbf{w}]|_{\mathbf{w}=\mathbf{z}}, \tag{17}$$

derived from a divergence $D[\mathbf{z} : \mathbf{w}]$. These two are dually coupled with respect to the Riemannian metric g_{ij} [1]. The meaning of dually coupled affine connections is not explained here, but will become clear in later sections, by using specific examples.

The Euclidean divergence, defined by

$$D_E[\mathbf{z} : \mathbf{w}] = \frac{1}{2} \sum (z_i - w_i)^2, \tag{18}$$

is a special case of divergence. We have

$$g_{ij} = \delta_{ij}, \tag{19}$$

$$\Gamma_{ijk} = \Gamma_{ijk}^* = 0, \tag{20}$$

where δ_{ij} is the Kronecker delta. Therefore, the derived geometry is Euclidean. Since it is self-dual ($\Gamma_{ijk} = \Gamma_{ijk}^*$), the duality does not play a role.

2.3 Invariant divergence: f -divergence

2.3.1 Information monotonicity

Let us consider a function $\mathbf{t}(\mathbf{x})$ of random variable \mathbf{x} , where \mathbf{t} and \mathbf{x} are vector-valued. We can derive probability distribution $\bar{p}(\mathbf{t}, \boldsymbol{\xi})$ of \mathbf{t} from $p(\mathbf{x}, \boldsymbol{\xi})$ by

$$\bar{p}(\mathbf{t}, \boldsymbol{\xi}) d\mathbf{t} = \int_{\mathbf{t}=\mathbf{t}(\mathbf{x})} p(\mathbf{x}, \boldsymbol{\xi}) d\mathbf{x}. \tag{21}$$

When $\mathbf{t}(\mathbf{x})$ is not reversible, that is, $\mathbf{t}(\mathbf{x})$ is a many-to-one mapping, there is loss of information by summarizing observed data \mathbf{x} into reduced $\mathbf{t} = \mathbf{t}(\mathbf{x})$. Hence, for two divergences between two distributions specified by $\boldsymbol{\xi}_1$ and $\boldsymbol{\xi}_2$,

$$D = D[p(\mathbf{x}, \boldsymbol{\xi}_1) : p(\mathbf{x}, \boldsymbol{\xi}_2)], \tag{22}$$

$$\bar{D} = D[\bar{p}(\mathbf{t}, \boldsymbol{\xi}_1) : \bar{p}(\mathbf{t}, \boldsymbol{\xi}_2)], \tag{23}$$

it is natural to require

$$\bar{D} \leq D. \tag{24}$$

The equality holds, when and only when $\mathbf{t}(\mathbf{x})$ is a sufficient statistics. A divergence function is said to be invariant, when this requirement is satisfied. The invariance is a key concept to construct information geometry of probability distributions [1,5].

Here, we use a simplified version of invariance due to Csiszár [18,19]. We consider the space S_n of all probability distributions over $n + 1$ atoms $X = \{x_0, x_1, \dots, x_n\}$. A probability distributions is given by $\mathbf{p} = (p_0, p_1, \dots, p_n)$, $p_i = \text{Prob}\{x = x_i\}$.

Let us divide X into m subsets, T_1, T_2, \dots, T_m ($m < n + 1$), say

$$T_1 = \{x_1, x_2, x_5\}, \quad T_2 = \{x_3, x_8, \dots\}, \quad \dots \quad (25)$$

This is a partition of X ,

$$X = \cup T_i, \quad (26)$$

$$T_i \cap T_j = \phi. \quad (27)$$

Let t be a mapping from X to $\{T_1, \dots, T_m\}$. Assume that we do not know the outcome x directly. Instead we can observe $t(x) = T_j$, knowing the subset T_j to which x belongs. This is called coarse-graining of X into $T = \{T_1, \dots, T_m\}$.

The coarse-graining generates a new probability distributions $\bar{p} = (\bar{p}_1, \dots, \bar{p}_m)$ over T_1, \dots, T_m ,

$$\bar{p}_j = \text{Prob} \{T_j\} = \sum_{x \in T_j} \text{Prob} \{x\}. \quad (28)$$

Let $\bar{D}[\bar{p} : \bar{q}]$ be an induced divergence between \bar{p} and \bar{q} . Since coarse-graining summarizes a number of elements into one subset, detailed information of the outcome is lost. Therefore, it is natural to require

$$\bar{D}[\bar{p} : \bar{q}] \leq D[\mathbf{p} : \mathbf{q}]. \quad (29)$$

When does the equality hold? Assume that the outcome x is known to belong to T_j . This gives some information to distinguish two distributions \mathbf{p} and \mathbf{q} . If we know further detail of x inside subset T_j , we obtain more information to distinguish the two probability distributions \mathbf{p} and \mathbf{q} . Since x belongs to T_j , we consider the two conditional probability distributions

$$p(x_i | x_i \in T_j), \quad q(x_i | x_i \in T_j) \quad (30)$$

under the condition that x is in subset T_j . If the two distributions are equal, we cannot obtain further information to distinguish \mathbf{p} from \mathbf{q} by observing x inside T_j . Hence,

$$\bar{D}[\bar{p} : \bar{q}] = D[\mathbf{p} : \mathbf{q}] \quad (31)$$

holds, when and only when

$$p(x_i | T_j) = q(x_i | T_j) \quad (32)$$

for all T_j and all $x_i \in T_j$, or

$$\frac{p_i}{q_i} = \lambda_j \quad (33)$$

for all $x_i \in T_j$ for some constant λ_j .

A divergence satisfying the above requirements is called an invariant divergence, and such a property is termed as information monotonicity.

2.3.2 f -divergence

The f -divergence was introduced by Csiszár [2] and also

by Ali and Silvey [20]. It is defined by

$$D_f[\mathbf{p} : \mathbf{q}] = \sum p_i f\left(\frac{q_i}{p_i}\right), \quad (34)$$

where f is a convex function satisfying

$$f(1) = 0. \quad (35)$$

For a function cf with a constant c , we have

$$D_{cf}[\mathbf{p} : \mathbf{q}] = cD_f[\mathbf{p} : \mathbf{q}]. \quad (36)$$

Hence, f and cf give the same divergence except for the scale factor c . In order to standardize the scale of divergence, we may assume that

$$f''(1) = 1, \quad (37)$$

provided f is differentiable. Further, for $f_c(u) = f(u) - c(u - 1)$ where c is any constant, we have

$$D_{f_c}[\mathbf{p} : \mathbf{q}] = D_f[\mathbf{p} : \mathbf{q}]. \quad (38)$$

Hence, we may use such an f that satisfies

$$f'(1) = 0 \quad (39)$$

without loss of generality. A convex function satisfying the above three conditions (35), (37), (39) is called a standard f function.

A divergence is said to be decomposable, when it is a sum of functions of components

$$D[\mathbf{p} : \mathbf{q}] = \sum_i D[p_i : q_i]. \quad (40)$$

The f -divergence (34) is a decomposable divergence.

Csiszár [18] found that any f -divergence satisfies information monotonicity. Moreover, the class of f -divergences is unique in the sense that any decomposable divergence satisfying information monotonicity is an f -divergence.

Theorem 1 Any f -divergence satisfies the information monotonicity. Conversely, any decomposable information monotonic divergence is written in the form of f -divergence.

A proof is found, e.g., in Ref. [21].

The Riemannian metric and affine connections derived from an f -divergence has a common invariant structure [1]. They are given by the Fisher information metric and $\pm\alpha$ -connections, which are shown in a later section.

An extensive list of f -divergences is given in Cichocki et al. [22]. Some of them are listed below.

- 1) The Kullback-Leibler (KL-) divergence: $f(u) = u \log u - (u - 1)$:

$$D_{\text{KL}}[\mathbf{p} : \mathbf{q}] = \sum p_i \log \frac{q_i}{p_i}. \quad (41)$$

2) Squared Hellinger distance: $f(u) = (\sqrt{u} - 1)^2$:

$$D_{\text{Hel}}[\mathbf{p} : \mathbf{q}] = \sum (\sqrt{p_i} - \sqrt{q_i})^2. \quad (42)$$

3) The α -divergence:

$$f_\alpha(u) = \frac{4}{1-\alpha^2} \left(1 - u^{\frac{1+\alpha}{2}}\right) - \frac{2}{1-\alpha}(u-1), \quad (43)$$

$$D_\alpha[\mathbf{p} : \mathbf{q}] = \frac{4}{1-\alpha^2} \sum \left(1 - p_i^{\frac{1+\alpha}{2}} q_i^{\frac{1-\alpha}{2}}\right). \quad (44)$$

The α -divergence was introduced by Havrda and Charvát [23], and has been studied extensively by Amari and Nagaoka [1]. Its applications were described earlier in Chernoff [24], and later in Matsuyama [25], Amari [26], etc. to mention a few. It is the squared Hellinger distance for $\alpha = 0$. The KL-divergence and its reciprocal are obtained in the limit of $\alpha \rightarrow \pm 1$. See Refs. [21,22,26] for the α -structure.

2.3.3 f -divergence of positive measures

We now extend the f -divergence to the space of positive measures M_n over $X = \{x_1, \dots, x_n\}$, whose points are given by the coordinates $\mathbf{z} = (z_1, \dots, z_n)$, $z_i > 0$. Here, z_i is the mass (measure) of x_i where the total mass $\sum z_i$ is positive and arbitrary. In many applications, \mathbf{z} is a non-negative array, and we can extend it to a non-negative double array $\mathbf{z} = (z_{ij})$, etc., that is, matrices and tensors. We first derive an f -divergence in M_n : For

two positive measures \mathbf{z} and \mathbf{y} , an f -divergence is given by

$$D_f[\mathbf{z} : \mathbf{y}] = \sum z_i f\left(\frac{y_i}{z_i}\right), \quad (45)$$

where f is a standard convex function. It should be noted that an f -divergence is no more invariant under the transformation from $f(u)$ to

$$f_c(u) = f(u) - c(u-1). \quad (46)$$

Hence, it is absolutely necessary to use a standard f in the case of M_n , because the conditions of divergence are violated otherwise.

Among all f -divergences, the α divergence [1,21,22] is given by

$$D_\alpha[\mathbf{z} : \mathbf{y}] = \sum z_i f_\alpha\left(\frac{y_i}{z_i}\right), \quad (47)$$

where

$$f_\alpha(u) = \begin{cases} \frac{4}{1-\alpha^2} \left(1 - u^{\frac{1+\alpha}{2}}\right) - \frac{2}{1-\alpha}(u-1), & \alpha \neq \pm 1, \\ u \log u - (u-1), & \alpha = 1, \\ -\log u + (u-1), & \alpha = -1, \end{cases} \quad (48)$$

and plays a special role in M_n . Here, we use a simple power function $u^{(1+\alpha)/2}$, changing it by adding a linear and a constant term to become the standard convex function $f_\alpha(u)$. It includes the logarithm as a limiting case.

The α -divergence is explicitly given in the following form [1,21,22]:

$$D_\alpha[\mathbf{z} : \mathbf{y}] = \begin{cases} \frac{2}{1-\alpha^2} \sum \left(\frac{1-\alpha}{2} z_i + \frac{1+\alpha}{2} y_i - z_i^{\frac{1+\alpha}{2}} y_i^{\frac{1-\alpha}{2}}\right), & \alpha \neq \pm 1, \\ \sum \left(z_i - y_i + y_i \log \frac{y_i}{z_i}\right), & \alpha = 1, \\ \sum \left(y_i - z_i + z_i \log \frac{z_i}{y_i}\right), & \alpha = -1. \end{cases} \quad (49)$$

2.4 Flat divergence: Bregman divergence

A divergence $D[\mathbf{z} : \mathbf{w}]$ gives a Riemannian metric g_{ij} by (13), and a pair of affine connections by (16), (17). When (20) holds, the manifold is flat. We study divergence functions which give flat structure in this subsection. In this case, coordinate system \mathbf{z} is flat, that is affine, although the metric $g_{ij}(\mathbf{z})$ is not Euclidean. When \mathbf{z} is affine, all the coordinate curves are regarded as straight lines, that is, geodesics. Here, we separate two concepts, flatness and metric. Mathematically speaking, we use an affine connection Γ_{ijk} to define flatness, which is not necessarily derived from the metric g_{ij} . Note that the Levi-Civita affine connection is derived from the met-

ric in Riemannian geometry, but our approach is more general.

2.4.1 Convex function

We show in this subsection that a dually flat geometry is derived from a divergence due to a convex function in terms of the Bregman divergence [3]. Conversely a dually flat manifold always has a convex function to define its geometry. The divergence derived from it is called the canonical divergence of a dually flat manifold [1].

Let $\psi(\mathbf{z})$ be a strictly convex differentiable function. Then, the linear hyperplane tangent to the graph

$$y = \psi(\mathbf{z}) \quad (50)$$

at \mathbf{z}_0 is

$$y = \nabla\psi(\mathbf{z}_0) \cdot (\mathbf{z} - \mathbf{z}_0) + \psi(\mathbf{z}_0), \quad (51)$$

and is always below the graph $y = \psi(\mathbf{z})$ (Fig. 4). Here, $\nabla\psi$ is the gradient,

$$\nabla\psi = \left(\frac{\partial}{\partial z_1}\psi(\mathbf{z}), \dots, \frac{\partial}{\partial z_n}\psi(\mathbf{z}) \right). \quad (52)$$

The difference of $\psi(\mathbf{z})$ and the tangent hyperplane is

$$D_\psi[\mathbf{z} : \mathbf{z}_0] = \psi(\mathbf{z}) - \psi(\mathbf{z}_0) - \nabla\psi(\mathbf{z}_0) \cdot (\mathbf{z} - \mathbf{z}_0) \geq 0. \quad (53)$$

See Fig. 4. This is called the Bregman divergence induced by $\psi(\mathbf{z})$.




Fig. 4 Bregman divergence due to convex function

This satisfies the conditions of a divergence. The induced Riemannian metric is

$$g_{ij}(\mathbf{z}) = \frac{\partial^2}{\partial z_i \partial z_j} \psi(\mathbf{z}), \quad (54)$$

and the coefficients of the affine connection vanish luckily,

$$\Gamma_{ijk}(\mathbf{z}) = 0. \quad (55)$$

Therefore, \mathbf{z} is an affine coordinate system, and a geodesic is always written as the linear form

$$\mathbf{z}(t) = \mathbf{a}t + \mathbf{b} \quad (56)$$

with parameter t and constant vectors \mathbf{a} and \mathbf{b} (see Ref. [27]).

2.4.2 Dual structure

In order to give a dual description related to a convex function $\psi(\mathbf{z})$, let us define its gradient,

$$\mathbf{z}^* = \nabla\psi(\mathbf{z}). \quad (57)$$

This is the Legendre transformation, and the correspondence between \mathbf{z} and \mathbf{z}^* is one-to-one. Hence, \mathbf{z}^* is regarded as another (nonlinear) coordinate system of S . In order to describe the geometry of S in terms of the dual coordinate system \mathbf{z}^* , we search for a dual convex

function $\psi^*(\mathbf{z}^*)$. We can obtain a dual potential function by

$$\psi^*(\mathbf{z}^*) = \max_{\mathbf{z}} \{ \mathbf{z} \cdot \mathbf{z}^* - \psi(\mathbf{z}) \}, \quad (58)$$

which is convex in \mathbf{z}^* . The two coordinate systems are dual, since the pair \mathbf{z} and \mathbf{z}^* satisfies the following relation:

$$\psi(\mathbf{z}) + \psi^*(\mathbf{z}^*) - \mathbf{z} \cdot \mathbf{z}^* = 0. \quad (59)$$

The inverse transformation is given by

$$\mathbf{z} = \nabla\psi^*(\mathbf{z}^*). \quad (60)$$

We can define the Bregman divergence $D_{\psi^*}[\mathbf{z}^* : \mathbf{w}^*]$ by using $\psi^*(\mathbf{z}^*)$. However, it is easy to prove

$$D_{\psi^*}[\mathbf{z}^* : \mathbf{w}^*] = D_\psi[\mathbf{w} : \mathbf{z}]. \quad (61)$$

Hence, they are essentially the same, that is the same when \mathbf{w} and \mathbf{z} are interchanged. It is enough to use only one common divergence.

By simple calculations, we see that the Bregman divergence can be rewritten in the dual form as

$$D[\mathbf{z} : \mathbf{w}] = \psi(\mathbf{z}) + \psi^*(\mathbf{w}^*) - \mathbf{z} \cdot \mathbf{w}^*. \quad (62)$$

We have thus two coordinate systems \mathbf{z} and \mathbf{z}^* of S . They define two flat structures such that a curve with parameter t ,

$$\mathbf{z}(t) = \mathbf{t}\mathbf{a} + \mathbf{b}, \quad (63)$$

is a ψ -geodesic, and

$$\mathbf{z}^*(t) = \mathbf{t}\mathbf{c}^* + \mathbf{d}^* \quad (64)$$

is a ψ^* -geodesic, where \mathbf{a} , \mathbf{b} , \mathbf{c}^* and \mathbf{d}^* are constants.

A Riemannian metric is defined by two metric tensors $G = (g_{ij})$ and $G^* = (g_{ij}^*)$,

$$g_{ij} = \frac{\partial^2}{\partial z_i \partial z_j} \psi(\mathbf{z}), \quad (65)$$

$$g_{ij}^* = \frac{\partial^2}{\partial z_i^* \partial z_j^*} \psi^*(\mathbf{z}^*), \quad (66)$$

in the two coordinate systems, and they are mutual inverses, $G^* = G^{-1}$ [1,27]. Because the squared local distance of two nearby points \mathbf{z} and $\mathbf{z} + d\mathbf{z}$ is given by

$$D[\mathbf{z} : \mathbf{z} + d\mathbf{z}] = \frac{1}{2} \sum g_{ij}(\mathbf{z}) dz_i dz_j, \quad (67)$$

$$= \frac{1}{2} \sum g_{ij}^*(\mathbf{z}^*) dz_i^* dz_j^*, \quad (68)$$

they give the same Riemannian distance. However, the two geodesics are different, which are also different from the Riemannian geodesic derived from the Riemannian metric. Hence, the minimality of the curve length does not hold for ψ - and ψ^* -geodesics.

Two geodesic curves (63) and (64) intersect at $t = 0$, when $\mathbf{b} = \mathbf{d}^*$. In such a case, they are orthogonal if their tangent vectors are orthogonal in the sense of the Riemannian metric. The orthogonality condition can be

represented in a simple form in terms of the two dual coordinates as

$$\langle \mathbf{a}, \mathbf{b}^* \rangle = \sum a_i b_i^* = 0. \quad (69)$$

2.4.3 Pythagorean theorem and projection theorem

A generalized Pythagorean theorem and projection theorem hold for a dually flat space [1]. They are highlights of a dually flat manifold.

Generalized Pythagorean Theorem. Let $\mathbf{r}, \mathbf{s}, \mathbf{q}$ be three points in S such that the ψ^* -geodesic connecting \mathbf{r} and \mathbf{s} is orthogonal to the ψ -geodesic connecting \mathbf{s} and \mathbf{q} (Fig. 5). Then,

$$D_\psi[\mathbf{r} : \mathbf{s}] + D_\psi[\mathbf{s} : \mathbf{q}] = D_\psi[\mathbf{r} : \mathbf{q}]. \quad (70)$$




Fig. 5 Pythagorean theorem

When S is a Euclidean space with

$$D[\mathbf{r} : \mathbf{s}] = \frac{1}{2} \sum (r_i - s_i)^2, \quad (71)$$

(70) is the Pythagorean theorem itself.

Now, we state the projection theorem. Let M be a submanifold of S (Fig. 6). Given $\mathbf{p} \in S$, a point $\mathbf{q} \in M$ is said to be the ψ -projection (ψ^* -projection) of \mathbf{p} to M when the ψ -geodesic (ψ^* -geodesic) connecting \mathbf{p} and \mathbf{q} is orthogonal to M with respect to the Riemannian metric g_{ij} .




Fig. 6 Projection of \mathbf{p} to M

Projection Theorem. Given $\mathbf{p} \in S$, the minimizer \mathbf{r}^* of $D_\psi(\mathbf{p}, \mathbf{r})$, $\mathbf{r} \in M$, is the ψ^* -projection of \mathbf{p} to M , and the minimizer \mathbf{r}^{**} of $D_\psi(\mathbf{r}, \mathbf{p})$, $\mathbf{r} \in M$, is the ψ -projection of \mathbf{p} to M .

The theorem is useful in various optimization problems searching for the closest point $\mathbf{r} \in M$ to a given \mathbf{p} .

2.4.4 Decomposable convex function and generalization

Let $U(z)$ be a convex function of a scalar z . Then, we have a simple convex function of \mathbf{z} ,

$$\psi(\mathbf{z}) = \sum U(z_i). \quad (72)$$

The dual of U is given by

$$U^*(z^*) = \max_z \{zz^* - U(z)\}, \quad (73)$$

and the dual coordinates are

$$z_i^* = U'(z_i). \quad (74)$$

The dual convex function is

$$\psi^*(\mathbf{z}^*) = \sum U^*(z_i^*). \quad (75)$$

The ψ -divergence due to U between \mathbf{z} and \mathbf{w} is decomposable and given by the sum of components [28–30],

$$D_\psi[\mathbf{z} : \mathbf{w}] = \sum \{U(z_i) + U^*(w_i^*) - z_i w_i^*\}. \quad (76)$$

We have discussed the flat structure given by a convex function due to $U(z)$. We now consider a more general case by using a nonlinear rescaling. Let us define by $r = k(z)$ a nonlinear scale r in the z -axis, where k is a monotone function. If we have a convex function $U(r)$ of the induced variable r , there emerges a new dually flat structure with a new convex function $U(r)$. Given two functions $k(z)$ and $U(r)$, we can introduce a new dually flat Riemannian structure in S , where the flat coordinates are not \mathbf{z} but $\mathbf{r} = k(\mathbf{z})$. There are infinitely many such structures depending on the choice of k and U . We show two typical examples, which give the α -divergence and β -divergence [22] in later sections. Both of them use a power function of the type u^q including the log and exponential function as the limiting cases.

2.5 Invariant flat structure of S

This section focuses on the manifolds which are equipped with both invariant and flat geometrical structure.

2.5.1 Exponential family of probability distributions

The following type of probability distributions of random variable \mathbf{y} parameterized by $\boldsymbol{\theta}$,

$$p(\mathbf{y}, \boldsymbol{\theta}) = \exp \left\{ \sum \theta_i s_i(\mathbf{y}) - \psi(\boldsymbol{\theta}) \right\}, \quad (77)$$

under a certain measure $\mu(\mathbf{y})$ on the space $Y = \{\mathbf{y}\}$ is called an exponential family. We denote it now by S . We introduce random variables $\mathbf{x} = (x_i)$,

$$x_i = s_i(\mathbf{y}), \tag{78}$$

and rewrite (77) in the standard form

$$p(\mathbf{x}, \boldsymbol{\theta}) = \exp \left\{ \sum \theta_i x_i - \psi(\boldsymbol{\theta}) \right\}, \tag{79}$$

where $\exp \{-\psi(\boldsymbol{\theta})\}$ is the normalization factor satisfying

$$\psi(\boldsymbol{\theta}) = \log \int \exp \left\{ \sum \theta_i x_i \right\} d\mu(\mathbf{x}). \tag{80}$$

This is called the free energy in the physics community and is a convex function.

We first give a few examples of exponential families.

1. Gaussian distributions S_G

A Gaussian distribution is given by

$$p(y; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left\{ -\frac{(y - \mu)^2}{2\sigma^2} \right\}. \tag{81}$$

Hence, all Gaussian distributions form a two-dimensional manifold S_G , where (μ, σ) plays the role of a coordinate system. Let us introduce new variables

$$x_1 = y, \tag{82}$$

$$x_2 = y^2, \tag{83}$$

$$\theta_1 = \frac{\mu}{2\sigma^2}, \tag{84}$$

$$\theta_2 = -\frac{1}{2\sigma^2}. \tag{85}$$

Then, (81) can be rewritten in the standard form

$$p(\mathbf{x}, \boldsymbol{\theta}) = \exp \{ \theta_1 x_1 + \theta_2 x_2 - \psi(\boldsymbol{\theta}) \}, \tag{86}$$

where

$$\psi(\boldsymbol{\theta}) = \frac{\mu^2}{2\sigma^2} - \log \left(\sqrt{2\pi}\sigma \right). \tag{87}$$

Here, $\boldsymbol{\theta} = (\theta_1, \theta_2)$ is a new coordinate system of S_G , called the natural or canonical coordinate system.

2. Discrete distributions S_n

Let y be a random variable taking values on $\{0, 1, \dots, n\}$. By defining

$$x_i = \delta_i(y), \quad i = 1, 2, \dots, n, \tag{88}$$

and

$$\theta_i = \log \frac{p_i}{p_0}, \tag{89}$$

(4) is rewritten in the standard form

$$p(\mathbf{x}, \boldsymbol{\theta}) = \exp \left\{ \sum \theta_i x_i - \psi(\boldsymbol{\theta}) \right\}, \tag{90}$$

where

$$\psi(\boldsymbol{\theta}) = -\log p_0. \tag{91}$$

From

$$p_0 = 1 - \sum p_i = 1 - p_0 \sum e^{\theta_i}, \tag{92}$$

we have

$$\psi(\boldsymbol{\theta}) = \log \left(1 + \sum e^{\theta_i} \right). \tag{93}$$

2.5.2 Geometry of exponential family

The Fisher information matrix is defined in a manifold of probability distributions by

$$g_{ij}(\boldsymbol{\theta}) = E \left[\frac{\partial}{\partial \theta_i} \log p(\mathbf{x}, \boldsymbol{\theta}) \frac{\partial}{\partial \theta_j} \log p(\mathbf{x}, \boldsymbol{\theta}) \right], \tag{94}$$

where E is expectation with respect to $p(\mathbf{x}, \boldsymbol{\theta})$. When S is an exponential family (79), it is calculated as

$$g_{ij}(\boldsymbol{\theta}) = \frac{\partial^2}{\partial \theta_i \partial \theta_j} \psi(\boldsymbol{\theta}). \tag{95}$$

This is a positive-definite matrix and $\psi(\boldsymbol{\theta})$ is a convex function. Hence, this is the Riemannian metric induced from a convex function ψ .

We can also prove that the invariant Riemannian metric derived from an f -divergence is exactly the same as the Fisher metric for any standard convex function f . It is an invariant metric.

A geometry is induced to $S = \{p(\mathbf{x}, \boldsymbol{\theta})\}$ through the convex function $\psi(\boldsymbol{\theta})$. Here, $\boldsymbol{\theta}$ is an affine coordinate system, which we call e -affine or e -flat coordinates, “ e ” representing the exponential structure of (79). By the Legendre transformation of ψ , we have the dual coordinates (we denote them by $\boldsymbol{\eta}$ instead of $\boldsymbol{\theta}^*$),

$$\boldsymbol{\eta} = \nabla \psi(\boldsymbol{\theta}). \tag{96}$$

By calculating the expectation of x_i , we have

$$E[x_i] = \eta_i = \frac{\partial}{\partial \theta_i} \psi(\boldsymbol{\theta}). \tag{97}$$

By this reason, the dual parameters $\boldsymbol{\eta} = (\eta_i)$ are called the expectation parameters of an exponential family. We call $\boldsymbol{\eta}$ the m -affine or m -flat coordinates, where “ m ” implying the mixture structure [1].

The dual convex function is given by the negative of entropy

$$\psi^*(\boldsymbol{\eta}) = \int p(\mathbf{x}, \boldsymbol{\theta}(\boldsymbol{\eta})) \log p(\mathbf{x}, \boldsymbol{\theta}(\boldsymbol{\eta})) d\mathbf{x}, \tag{98}$$

where $p(\mathbf{x}, \boldsymbol{\theta}(\boldsymbol{\eta}))$ is considered as a function of $\boldsymbol{\eta}$. Then, the Bregman divergence is

$$D[\boldsymbol{\theta}_1 : \boldsymbol{\theta}_2] = \psi(\boldsymbol{\theta}_1) + \psi^*(\boldsymbol{\eta}_2) - \boldsymbol{\theta}_1 \cdot \boldsymbol{\eta}_2, \tag{99}$$

where $\boldsymbol{\eta}_i$ is the dual coordinates of $\boldsymbol{\theta}_i$, $i = 1, 2$.

We show that manifolds of exponential families have both invariant and flat geometrical structure. However, they are not unique and it is possible to introduce both invariant flat structure to a mixture family of probability distributions.

Theorem 2 The Bregman divergence is the Kullback-Leibler divergence given by

$$\text{KL} [p(\mathbf{x}, \boldsymbol{\theta}_1) : p(\mathbf{x}, \boldsymbol{\theta}_2)] = \int p(\mathbf{x}, \boldsymbol{\theta}_1) \log \frac{p(\mathbf{x}, \boldsymbol{\theta}_1)}{p(\mathbf{x}, \boldsymbol{\theta}_2)} d\mathbf{x}. \tag{100}$$

This shows that the KL-divergence is the canonical divergence derived from the dual flat structure of $\psi(\boldsymbol{\theta})$. Moreover, it is proved that the KL-divergence is the unique divergence belonging to both the f -divergence and Bregman divergence classes for manifolds of probability distributions.

For a manifold of positive measures, there are divergences other than the KL-divergence, which are invariant and dually flat, that is, belonging to the classes f -divergences and Bregman divergences at the same time. See the next subsection.

2.5.3 Invariant and flat structure in manifold of positive measures

We consider a manifold M of positive arrays \mathbf{z} , $z_i > 0$, where $\sum z_i$ can be arbitrary. We define an invariant and flat structure in M . Here, \mathbf{z} is a coordinate system, and consider the f -divergence given by

$$D_f[\mathbf{z} : \mathbf{w}] = \sum_{i=1}^n z_i f\left(\frac{w_i}{z_i}\right). \tag{101}$$

We now introduce a new coordinate system defined by

$$r_i^{(\alpha)} = k_\alpha(z_i), \tag{102}$$

where the α -representation of z_i is given by

$$k_\alpha(u) = \begin{cases} \frac{2}{1-\alpha} \left(u^{\frac{1-\alpha}{2}} - 1\right), & \alpha \neq 1, \\ \log u, & \alpha = 1. \end{cases} \tag{103}$$

Then, the α -coordinates \mathbf{r}_α of M are defined by

$$r_i^{(\alpha)} = k_\alpha(z_i). \tag{104}$$

We use a convex function $U_\alpha(r)$ of r defined by

$$U_\alpha(r) = \frac{2}{1+\alpha} k_\alpha^{-1}(r). \tag{105}$$

In terms of z , this is a linear function,

$$U_\alpha\{r(z)\} = \frac{2}{1+\alpha} z, \tag{106}$$

which is not a (strictly) convex function of z but is convex with respect to r . The α -potential function defined by

$$\begin{aligned} \psi_\alpha(\mathbf{r}) &= \frac{2}{1+\alpha} \sum k_\alpha^{-1}(r_i) \\ &= \frac{2}{1+\alpha} \sum \left(1 + \frac{1-\alpha}{2} r_i\right)^{\frac{2}{1-\alpha}} \end{aligned} \tag{107}$$

is a convex function of \mathbf{r} .

The dual potential is simply given by

$$\psi_\alpha^*(\mathbf{r}^*) = \psi_{-\alpha}(\mathbf{r}^*), \tag{108}$$

and the dual affine coordinates are

$$\mathbf{r}^{*(\alpha)} = \mathbf{r}^{(-\alpha)} = k_{-\alpha}(\mathbf{z}). \tag{109}$$

We can then prove the following theorem.

Theorem 3 The α -divergence of M is the Bregman divergence derived from the α -representation k_α and the linear function U_α of z .

Proof We have the Bregman divergence between $\mathbf{r} = k(\mathbf{z})$ and $\mathbf{s} = k(\mathbf{y})$ based on ψ_α as

$$D_\alpha[\mathbf{r} : \mathbf{s}] = \psi_\alpha(\mathbf{r}) + \psi_{-\alpha}(\mathbf{s}^*) - \mathbf{r} \cdot \mathbf{s}^*, \tag{110}$$

where \mathbf{s}^* is the dual of \mathbf{s} . By substituting (107), (108) and (109) in (110), we see that $D_\alpha[\mathbf{r}(\mathbf{z}) : \mathbf{s}(\mathbf{y})]$ is equal to the α -divergence defined in (49).

This proves that the α -divergence for any α belongs to the intersection of the classes of f -divergences and Bregman divergences in M . Hence, it possesses information monotonicity and the induced geometry is dually flat at the same time. However, the constraint $\sum z_i = 1$ is not linear in the flat coordinates \mathbf{r}^α or $\mathbf{r}^{-\alpha}$ except for the case $\alpha = \pm 1$. Hence, although M is flat for any α , the manifold P of probability distributions is not dually flat except for $\alpha = \pm 1$, and it is a curved submanifold in M . Hence, the α -divergences ($\alpha \neq \pm 1$) do not belong to the class of Bregman divergences in P . This proves that the KL-divergence belongs to both classes of divergences in P and is unique.

The α -divergences are used in many applications, e.g., Refs. [21,25,26].

Theorem 4 The α -divergence is the unique class of divergences sitting at the intersection of the f -divergence and Bregman divergence classes.

Proof The f -divergence is of the form (45) and the decomposable Bregman divergence is of the form

$$D[\mathbf{z} : \mathbf{y}] = \sum \tilde{U}(z_i) + \sum \tilde{U}^*(y_i) - \sum r_i(z_i) r_i^*(y_i), \tag{111}$$

where

$$\tilde{U}(z) = U(k(z)), \tag{112}$$

and so on. When they are equal, we have f , r , and r^* that satisfy

$$zf\left(\frac{y}{z}\right) = r(z)r^*(y) \tag{113}$$

except for additive terms depending only on z or y . Differentiating the above by y , we get

$$f'\left(\frac{y}{z}\right) = r(z)r^{*'}(y). \tag{114}$$

By putting $x = y, y = 1/z$, we get

$$f'(xy) = r\left(\frac{1}{y}\right)r^{*'}(x). \tag{115}$$

Hence, by putting $h(u) = \log f'(u)$, we have

$$h(xy) = s(x) + t(y), \tag{116}$$

where $s(x) = \log r^{*'}(x), t(y) = r(1/y)$. By differentiating the above equation with respect to x and putting $x = 1$, we get

$$h'(y) = \frac{c}{y}. \tag{117}$$

This proves that f is of the form

$$f(u) = \begin{cases} cu^{\frac{1+\alpha}{2}}, & \alpha \neq \pm 1, \\ cu \log u, & \alpha = 1, \\ c \log u, & \alpha = -1. \end{cases} \tag{118}$$

By changing the above into the standard form, we arrive at the α -divergence.

2.6 β -divergence

It is useful to introduce a family of β -divergences, which are not invariant but dually flat, having the structure of Bregman divergences. This is used in the field of machine learning and statistics [28,29].

Use \mathbf{z} itself as a coordinate system of M ; that is, the representation function is $k(z) = z$. The β -divergence [30] is induced by the potential function

$$U_\beta(z) = \begin{cases} \frac{1}{\beta+1}(1+\beta z)^{\frac{\beta+1}{\beta}}, & \beta > 0, \\ \exp z, & \beta = 0. \end{cases} \tag{119}$$

The β -divergence is thus written as

$$D_\beta[\mathbf{z} : \mathbf{y}] = \begin{cases} \frac{1}{\beta+1} \sum (y_i^{\beta+1} - z_i^{\beta+1}) - \frac{1}{\beta} \sum z_i^\beta (y_i - z_i), & \beta > 0, \\ \sum \left(z_i \log \frac{z_i}{y_i} + y_i - z_i \right), & \beta = 0. \end{cases} \tag{120}$$

It is the KL-divergence when $\beta = 0$, but it is different from the α -divergence when $\beta > 0$.

Minami and Eguchi [30] demonstrated that statistical inference based on the β -divergence ($\beta > 0$) is robust. Such idea has been applied to machine learning in Murata et al. [29]. The β -divergence induces a dually flat structure in M . Since its flat coordinates are \mathbf{z} , the restriction

$$\sum z_i = 1 \tag{121}$$

is a linear constraint. Hence, the manifold P of the probability distributions is also dually flat, where $\mathbf{z}, \sum z_i = 1$, is its flat coordinates. The dual flat coordinates are

$$z_i^* = \begin{cases} (1 + \beta z_i)^{\frac{1}{\beta}}, & \beta \neq 0, \\ \exp z_i, & \beta = 0, \end{cases} \tag{122}$$

depending on β .

3 Information geometry of alternative minimization

Given a divergence function $D[\mathbf{z} : \mathbf{w}]$ on a space S , we encounter the problem of minimizing $D[\mathbf{z} : \mathbf{w}]$ under the constraints that \mathbf{z} belongs to a submanifold $M \subset S$ and \mathbf{w} belongs to another submanifold $E \subset S$. A typical example is the EM algorithm [6], where M represents a manifold of partially observed data and E represents

a statistical model to which the true distribution is assumed to belong [8]. The EM algorithm is popular in statistics, machine learning and others [7,31,32].

3.1 Geometry of alternative minimization

When a divergence function $D[\mathbf{z} : \mathbf{w}]$ is given in a manifold S , we consider two submanifolds M and E , where we assume $\mathbf{z} \in M$ and $\mathbf{w} \in E$. The divergence of two submanifolds M and E is given by

$$D[M : E] = \min_{\mathbf{z} \in M, \mathbf{w} \in E} D[\mathbf{z} : \mathbf{w}] = D[\mathbf{z}^* : \mathbf{w}^*], \tag{123}$$

where $\mathbf{z}^* \in M$ and $\mathbf{w}^* \in E$ are a pair of the points that attain the minimum. When M and E intersect, $D[M : E] = 0$, and $\mathbf{z}^* = \mathbf{w}^*$ lies at the intersection.

An alternative minimization is a procedure to calculate $(\mathbf{z}^*, \mathbf{w}^*)$ iteratively (see Fig. 7):

- 1) Take an initial point $\mathbf{z}^0 \in M$ for $t = 0$. Repeat steps 2 and 3 for $t = 0, 1, 2, \dots$ until convergence.
- 2) Calculate

$$\mathbf{w}^t = \arg \min_{\mathbf{w} \in E} D[\mathbf{z}^t : \mathbf{w}]. \tag{124}$$

- 3) Calculate

$$\mathbf{z}^{t+1} = \arg \min_{\mathbf{z} \in M} D[\mathbf{z} : \mathbf{w}^{t+1}]. \tag{125}$$

When the direct minimization with respect to one




Fig. 7 Alternative minimization

variable is computationally difficult, we may use a decremental procedure:

$$\mathbf{w}^{t+1} = \mathbf{w}^t - \eta_t \nabla_{\mathbf{w}} D[\mathbf{z}^t : \mathbf{w}^t], \quad (126)$$

$$\mathbf{z}^{t+1} = \mathbf{z}^t - \eta_t \nabla_{\mathbf{z}} D[\mathbf{z}^{t+1} : \mathbf{w}^t], \quad (127)$$

where η_t is a learning constant and $\nabla_{\mathbf{w}}$ and $\nabla_{\mathbf{z}}$ are gradients with respect to \mathbf{w} and \mathbf{z} , respectively.

It should also be noted that the natural gradient (Riemannian gradient of D) [33] is given by

$$\tilde{\nabla}_{\mathbf{z}} D[\mathbf{z} : \mathbf{w}] = G^{-1}(\mathbf{z}) \frac{\partial}{\partial \mathbf{z}} D[\mathbf{z} : \mathbf{w}], \quad (128)$$

where

$$G = \frac{\partial^2}{\partial \mathbf{z} \partial \mathbf{z}} D[\mathbf{z} : \mathbf{w}]|_{\mathbf{w}=\mathbf{z}}. \quad (129)$$

Hence, this becomes the Newton-type method.

3.2 Alternative projections

When $D[\mathbf{z} : \mathbf{w}]$ is a Bregman divergence derived from a convex function ψ , the space S is equipped with a dually flat structure. Given \mathbf{w} , the minimization of $D[\mathbf{z} : \mathbf{w}]$ with respect to $\mathbf{z} \in M$ is given by the ψ -projection of \mathbf{w} to M ,

$$\arg \min_{\mathbf{z}} D[\mathbf{z} : \mathbf{w}] = \prod_M^{(\psi)} \mathbf{w}. \quad (130)$$

This is unique when M is ψ^* -flat. On the other hand, given \mathbf{z} , the minimization with respect to \mathbf{w} is given by the ψ^* -projection,

$$\arg \min_{\mathbf{w}} D[\mathbf{z} : \mathbf{w}] = \prod_M^{(\psi^*)} \mathbf{z}. \quad (131)$$

This is unique when M is ψ -flat.

3.3 EM algorithm

Let us consider a statistical model $p(\mathbf{z}; \boldsymbol{\xi})$ with random variables \mathbf{z} , parameterized by $\boldsymbol{\xi}$. The set of all such distributions $E = \{p(\mathbf{z}; \boldsymbol{\xi})\}$ forms a submanifold included in the entire manifold of probability distributions

$$S = \{p(\mathbf{z})\}. \quad (132)$$

Assume that vector random variable \mathbf{z} is divided into two parts $\mathbf{z} = (\mathbf{h}, \mathbf{x})$ and the values of \mathbf{x} are observed but \mathbf{h} are not. Non-observed random variables $\mathbf{h} = (h_i)$ are called missing data or hidden variables. It is possible to eliminate unobservable \mathbf{h} by summing up $p(\mathbf{h}, \mathbf{x})$ over all \mathbf{h} , and we have a new statistical model,

$$\tilde{p}(\mathbf{x}, \boldsymbol{\xi}) = \sum_{\mathbf{h}} p(\mathbf{h}, \mathbf{x}, \boldsymbol{\xi}), \quad (133)$$

$$\tilde{E} = \{\tilde{p}(\mathbf{x}; \boldsymbol{\xi})\}. \quad (134)$$

This $\tilde{p}(\mathbf{x}, \boldsymbol{\xi})$ is the marginal distribution of $p(\mathbf{h}, \mathbf{x}; \boldsymbol{\xi})$. Based on observed \mathbf{x} , we can estimate $\hat{\boldsymbol{\xi}}$, for example, by maximizing the likelihood,

$$\hat{\boldsymbol{\xi}} = \arg \max_{\boldsymbol{\xi}} \tilde{p}(\mathbf{x}, \boldsymbol{\xi}). \quad (135)$$

However, in many problems $p(\mathbf{z}; \boldsymbol{\xi})$ has a simple form while $\tilde{p}(\mathbf{x}; \boldsymbol{\xi})$ does not. The marginal distribution $\tilde{p}(\mathbf{x}, \boldsymbol{\xi})$ has a complicated form and it is computationally difficult to estimate $\boldsymbol{\xi}$ from \mathbf{x} . It is much easier to calculate $\hat{\boldsymbol{\xi}}$ when \mathbf{z} is observed. The EM algorithm is a powerful method used in such a case [6].

The EM algorithm consists of iterative procedures to use values of the hidden variables \mathbf{h} guessed from observed \mathbf{x} and the current estimator $\hat{\boldsymbol{\xi}}$. Let us assume that n iid data $\mathbf{z}_t = (\mathbf{h}_t, \mathbf{x}_t)$, $t = 1, \dots, n$, are generated from a distribution $p(\mathbf{z}, \boldsymbol{\xi})$. If data \mathbf{h}_t are not missing, we have the following empirical distribution:

$$\bar{p}(\mathbf{h}, \mathbf{x} | \bar{\mathbf{h}}, \bar{\mathbf{x}}) = \frac{1}{n} \sum_t \delta(\mathbf{x} - \mathbf{x}_t) \delta(\mathbf{h} - \mathbf{h}_t), \quad (136)$$

where $\bar{\mathbf{h}}$ and $\bar{\mathbf{x}}$ represent the sets of data $(\mathbf{h}_1, \mathbf{h}_2, \dots)$ and $(\mathbf{x}_1, \mathbf{x}_2, \dots)$ and δ is the delta function. When S is an exponential family, $\bar{p}(\mathbf{h}, \mathbf{x} | \bar{\mathbf{h}}, \bar{\mathbf{x}})$ is the distribution given by the sufficient statistics composed of $(\bar{\mathbf{h}}, \bar{\mathbf{x}})$.

The hidden variables are not observed in reality, and the observed part $\bar{\mathbf{x}}$ cannot determine the empirical distribution $\bar{p}(\mathbf{x}, \mathbf{y})$ (136). To overcome this difficulty, let us use a conditional probability $q(\mathbf{h}|\mathbf{x})$ of \mathbf{h} when \mathbf{x} is observed. We then have

$$p(\mathbf{h}, \mathbf{x}) = q(\mathbf{h}|\mathbf{x})p(\mathbf{x}). \quad (137)$$

In the present case, the observed part $\bar{\mathbf{x}}$ determines the empirical distribution $\bar{p}(\mathbf{x})$,

$$\bar{p}(\mathbf{x}) = \frac{1}{n} \sum_t \delta(\mathbf{x} - \mathbf{x}_t), \quad (138)$$

but the conditional part $q(\mathbf{h}|\mathbf{x})$ remains unknown. Taking this into account, we define a submanifold based on the observed data $\bar{\mathbf{x}}$,

$$M(\bar{\mathbf{x}}) = \{p(\mathbf{h}, \mathbf{x}) | p(\mathbf{h}, \mathbf{x}) = q(\mathbf{h}|\mathbf{x})\bar{p}(\mathbf{x})\}, \quad (139)$$

where $q(\mathbf{h}|\mathbf{x})$ is arbitrary.

In the observable case, we have an empirical distribution $\bar{p}(\mathbf{h}, \mathbf{x}) \in S$ from the observation, which does not usually belong to our model E . The maximum likelihood estimator $\hat{\xi}$ is the point in E that minimizes the KL-divergence from \bar{p} to E ,

$$\hat{\xi} = \arg \min \text{KL} [\bar{p}(\mathbf{h}, \mathbf{x}) : p(\mathbf{x}, \xi)]. \quad (140)$$

In the present case, we do not know \bar{p} but know $M(\bar{\mathbf{x}})$ which includes the unknown $\bar{p}(\mathbf{h}, \mathbf{x})$. Hence, we search for $\xi \in E$ and $\hat{q}(\mathbf{h}|\mathbf{x}) \in M(\bar{\mathbf{x}})$ that minimizes

$$\text{KL}[M : E]. \quad (141)$$

The pair of minimizers gives the points which define the divergence of two submanifolds $M(\bar{\mathbf{x}})$ and E . The maximum likelihood estimator is a minimizer of (141).

Given partial observations $\bar{\mathbf{x}}$, the EM algorithm consists of the following interactive procedures, $t = 0, 1, \dots$. We begin with an arbitrary initial guess ξ_0 and $q_0(\mathbf{h}|\mathbf{x})$, and construct a candidate $p_t(\mathbf{h}, \mathbf{x}) = q_t(\mathbf{h}|\mathbf{x})\bar{p}(\mathbf{x}) \in M = M(\bar{\mathbf{x}})$ for $t = 0, 1, 2, \dots$. We search for the next guess ξ_{t+1} that minimizes $\text{KL}[p_t(\mathbf{h}, \mathbf{x}) : p(\mathbf{h}, \mathbf{x}, \xi)]$. This is the same as maximizing the pseudo-log-likelihood

$$L_t(\xi) = \sum_{\mathbf{h}, \mathbf{x}_i} p_t(\mathbf{h}, \mathbf{x}_i) \log p(\mathbf{h}, \mathbf{x}_i; \xi). \quad (142)$$

This is called the M -step, that is the m -projection of $p_t(\mathbf{h}, \mathbf{x}) \in M$ to E . Let the maximum value be ξ_{t+1} , which is the estimator at $t + 1$.

We then search for the next candidate $p_t(\mathbf{h}, \mathbf{x}) \in M$ that minimizes $\text{KL}[p(\mathbf{h}, \mathbf{x}) : p(\mathbf{h}, \mathbf{x}, \xi_{t+1})]$. The minimizer is denoted by $p_{t+1}(\mathbf{h}, \mathbf{x})$. This distribution is used to calculate the expectation of the new log likelihood function $L_{t+1}(\xi)$. Hence, this is called the E -step. The following theorem [8] plays a fundamental role in calculating the expectation with respect to $p_t(\mathbf{h}, \mathbf{x})$.

Theorem 5 Along the e -geodesic of projecting $p(\mathbf{h}, \mathbf{x}; \xi_t)$ to $M(\bar{\mathbf{x}})$, the conditional probability distribution $q_t(\mathbf{h}|\mathbf{x})$ is kept invariant. Hence, by the e -projection of $p(\mathbf{h}, \mathbf{x}; \xi_t)$ to M , the resultant conditional distribution is given by

$$q_t(\mathbf{h}|\mathbf{x}_i) = p(\mathbf{h}|\mathbf{x}_i; \xi_t). \quad (143)$$

This shows that the e -projection is given by

$$q_t(\mathbf{h}, \mathbf{x}) = p(\mathbf{h}|\mathbf{x}; \xi_t) \delta(\mathbf{x} - \bar{\mathbf{x}}). \quad (144)$$

Hence, the theorem implies that the next expected log likelihood is calculated by

$$L_{t+1}(\xi) = \sum_{\mathbf{h}, \mathbf{x}_i} p(\mathbf{h}|\mathbf{x}_i; \xi_{t+1}) \log p(\mathbf{h}, \mathbf{x}_i; \xi), \quad (145)$$

which is to be maximized.

The EM algorithm is formally described as follows.

- 1) Begin with an arbitrary initial guess ξ_0 at $t = 0$, and repeat the E -step and M -step for $t = 0, 1, 2, \dots$, until convergence.
- 2) E -step: Use the e -projection of $p(\mathbf{h}, \mathbf{x}; \xi_t)$ to $M(\bar{\mathbf{x}})$ to obtain

$$q_t(\mathbf{h}|\mathbf{x}) = p(\mathbf{h}|\mathbf{x}; \xi_t), \quad (146)$$

and calculate the conditional expectation

$$L_t(\xi) = \sum_{\mathbf{h}, \mathbf{x}_i} p(\mathbf{h}|\mathbf{x}_i; \xi_t) \log p(\mathbf{h}, \mathbf{x}_i; \xi). \quad (147)$$

- 3) M -step: Calculate the maximizer of $L_t(\xi)$, that is the m -projection of $p_t(\mathbf{h}, \mathbf{x})$ to E ,

$$\xi_{t+1} = \arg \max L_t(\xi). \quad (148)$$

4 Information geometry of Ying-Yang machine

4.1 Recognition model and generating model

We study the Ying-Yang mechanism proposed by Xu [9–14] from the information geometry point of view. Let us consider a hierarchical stochastic system, which consists of two vector random variables \mathbf{x} and \mathbf{y} . Let \mathbf{x} be a lower level random variable representing a primitive description of the world, while let \mathbf{y} be a higher level random variable representing a more conceptual or elaborated description. Obviously, \mathbf{x} and \mathbf{y} are closely correlated. One may consider them as information representations in the brain. Sensory information activates neurons in a primitive sensory area of the brain, and its firing pattern is represented by \mathbf{x} . Components x_i of \mathbf{x} can be binary variables taking values 0 and 1, or analog values representing firing rates of neurons. The conceptual information activates a higher-order area of the brain, with a firing pattern \mathbf{y} .

A probability distribution $p(\mathbf{x})$ of \mathbf{x} reflects the information structure of the outer world. When \mathbf{x} is observed, the brain processes this primitive information and activates a higher-order area to recognize its higher-order representation \mathbf{y} . Its mechanism is represented by a conditional probability distribution $p(\mathbf{y}|\mathbf{x})$. This is possibly specified by a parameter ξ , in the form of $p(\mathbf{y}|\mathbf{x}; \xi)$. Then, the joint distribution is given by

$$p(\mathbf{y}, \mathbf{x}; \xi) = p(\mathbf{y}|\mathbf{x}; \xi)p(\mathbf{x}). \quad (149)$$

The parameters ξ would be synaptic weights of neurons to generate pattern \mathbf{y} . When we receive information \mathbf{x} , it is transformed to higher-order conceptual information \mathbf{y} . This is the bottom-up process and we call this a recognition model.

Our brain can work in the reverse way. From conceptual information pattern \mathbf{y} , a primitive pattern \mathbf{x} will

be generated by the top-down neural mechanism. It is represented by a conditional probability $q(\mathbf{x}|\mathbf{y})$, which will be parameterized by ζ , as $q(\mathbf{x}|\mathbf{y}; \zeta)$. Here, ζ will represent the top-down neural mechanism. When the probability distribution of \mathbf{y} is $q(\mathbf{y}; \zeta)$, the entire process is represented by another joint probability distribution

$$q(\mathbf{y}, \mathbf{x}; \zeta) = q(\mathbf{y}; \zeta)q(\mathbf{x}|\mathbf{y}; \zeta). \quad (150)$$

We have thus two stochastic mechanisms $p(\mathbf{y}, \mathbf{x}; \xi)$ and $q(\mathbf{y}, \mathbf{x}; \zeta)$. Let us denote by S the set of all the joint probability distributions of \mathbf{x} and \mathbf{y} . The recognition model

$$M_R = \{p(\mathbf{y}, \mathbf{x}; \xi)\} \quad (151)$$

forms its submanifold, and the generative model

$$M_G = \{q(\mathbf{y}, \mathbf{x}; \zeta)\} \quad (152)$$

forms another submanifold.

This type of information mechanism is proposed by Xu [9,14], and named the Ying-Yang machine. The Yang machine (male machine) is responsible for the recognition model $\{p(\mathbf{y}, \mathbf{x}; \xi)\}$ and the Ying machine (female machine) is responsible for the generative model $\{q(\mathbf{y}, \mathbf{x}; \zeta)\}$. These two models are different, but should be put in harmony (Ying-Yang harmony).

4.2 Ying-Yang harmony

When a divergence function $D[p(\mathbf{x}, \mathbf{y}) : q(\mathbf{x}, \mathbf{y})]$ is defined in S , we have the divergence between the two models M_R and M_G ,

$$D[M_R : M_G] = \min_{\xi, \zeta} D[p(\mathbf{y}, \mathbf{x}; \xi) : q(\mathbf{y}, \mathbf{x}; \zeta)]. \quad (153)$$

When M_R and M_G intersect,

$$D[M_R : M_G] = 0 \quad (154)$$

at the intersecting points. However, they do not intersect in general,

$$M_R \cap M_G = \phi. \quad (155)$$

It is desirable that the recognition and generating models are closely related. The minimizer of D gives such a harmonized pair called the Ying-Yang harmony [14].

A typical divergence is the Kullback-Leibler divergence $\text{KL}[p : q]$, which has merits of being invariant and generating dually flat geometrical structure. In this case, we can define the two projections, e -projection and m -projection. Let $p^*(\mathbf{y}, \mathbf{x})$ and $q^*(\mathbf{y}, \mathbf{x})$ be the pair of minimizers of $D[p : q]$. Then, the m -projection of q^* to M_R is p^* ,

$$\prod_{M_R}^{(m)} q^* = p^*, \quad (156)$$

and the e -projection of p^* to M_G is q^* ,

$$\prod_{M_G}^{(e)} p^* = q^*. \quad (157)$$

These relations hold for any dually flat divergence, generated by a convex function ψ over S .

An iterative algorithm to realize the Ying-Yang harmony is, for $t = 0, 1, 2, \dots$,

$$p_{t+1} = \prod_{M_R}^{(m)} q_t, \quad (158)$$

$$q_{t+1} = \prod_{M_G}^{(e)} p_{t+1}. \quad (159)$$

This is a typical alternative minimization algorithm.

4.3 Various models of Ying-Yang structure

A Yang machine receives signal \mathbf{x} from the outside, and processes it by generating a higher-order signal \mathbf{y} , which is stochastically generated by using the conditional probability $p(\mathbf{y}|\mathbf{x}, \xi)$. The performance of the machine is improved by learning, where the parameter ξ is modified step-by-step. We give two typical examples of information processing; pattern classification and self-organization.

Learning of a pattern classifier is a supervised learning process, where teacher signal z is given.

1. pattern classifier

Let $C = \{C_1, \dots, C_m\}$ denote the set of categories, and each \mathbf{x} belongs to one of them. Given an \mathbf{x} , a machine processes it and gives an answer \mathbf{y}_κ , where \mathbf{y}_κ , $\kappa = 1, \dots, m$, correspond to one of the categories. In other words, \mathbf{y}_κ is a pattern representing C_κ . The output \mathbf{y} is generated from \mathbf{x} :

$$\mathbf{y} = \mathbf{f}(\mathbf{x}, \xi) + \mathbf{n}, \quad (160)$$

where \mathbf{n} is a zero-mean Gaussian noise with covariance matrix $\sigma^2 I$, I being the identity matrix. Then, the conditional distribution is given by

$$p(\mathbf{y}|\mathbf{x}; \xi) = c \exp \left\{ -\frac{1}{2\sigma^2} |\mathbf{y} - \mathbf{f}(\mathbf{x}, \xi)|^2 \right\}, \quad (161)$$

where c is the normalization constant.

The function $\mathbf{f}(\mathbf{x}, \xi)$ is specified by a number of parameters ξ . In the case of multilayer perceptron, it is given by

$$\mathbf{y}_i = \sum_j v_{ij} \varphi \left\{ \sum_k w_{jk} x_k - h_j \right\}. \quad (162)$$

Here, w_{jk} is the synaptic weight from input x_k to the j th hidden unit, h_j is its threshold, and φ is a sigmoidal function. The output of the j th hidden unit is hence $\varphi(\sum_k w_{jk} x_k - h_j)$ and is connected linearly to the i th output unit with synaptic weight v_{ij} .

In the case of supervised learning, a teacher signal \mathbf{y}_κ representing the category C_κ is given, when the input

pattern is $\mathbf{x} \in C_\kappa$. Since the current output of the Yang machine is $\mathbf{f}(\mathbf{x}, \boldsymbol{\xi})$, the error of the machine is

$$l(\boldsymbol{\xi}, \mathbf{x}, \mathbf{y}_\kappa) = \frac{1}{2} |\mathbf{y}_\kappa - \mathbf{f}(\mathbf{x}, \boldsymbol{\xi})|^2. \quad (163)$$

Therefore, the stochastic gradient learning rule modifies the current $\boldsymbol{\xi}$ to $\boldsymbol{\xi} + \Delta\boldsymbol{\xi}$ by

$$\Delta\boldsymbol{\xi} = -\varepsilon \frac{\partial l}{\partial \boldsymbol{\xi}}, \quad (164)$$

where ε is the learning constant and $\nabla l = \left(\frac{\partial l}{\partial \xi_1}, \dots, \frac{\partial l}{\partial \xi_m} \right)$ is the gradient of l with respect to $\boldsymbol{\xi}$. Since the space of $\boldsymbol{\xi}$ is Riemannian, the natural gradient learning rule [33] is given by

$$\Delta\boldsymbol{\xi} = -\varepsilon G^{-1} \nabla l, \quad (165)$$

where G is the Fisher information matrix derived from the conditional distribution.

2. Non-supervised learning and clustering

When signals \mathbf{x} in the outside world are divided into a number of clusters, it is expected that Yang machine reproduces such a clustering structure. An example of clustered signals is the Gaussian mixture

$$p(\mathbf{x}) = \sum_{\kappa=1}^m w_\kappa \exp \left\{ -\frac{1}{2\sigma^2} |\mathbf{x} - \mathbf{y}_\kappa|^2 \right\}, \quad (166)$$

where \mathbf{y}_κ is the center of cluster C_κ . Given \mathbf{x} , one can see which cluster it belongs to by calculating the output $\mathbf{y} = \mathbf{f}(\mathbf{x}, \boldsymbol{\xi}) + \mathbf{n}$. Here, \mathbf{y} is an inner representation of the clusters and its calculation is controlled by $\boldsymbol{\xi}$. There are a number of clustering algorithms. See, e.g., Ref. [9]. We can use them to modify $\boldsymbol{\xi}$. Since no teacher signals exist, this is unsupervised learning.

Another typical non-supervised learning is implemented by the self-organization model given by Amari and Takeuchi [34].

The role of the Ying machine is to generate a pattern $\tilde{\mathbf{x}}$ from a conceptual information pattern \mathbf{y} . Since the mechanism of generating $\tilde{\mathbf{x}}$ is governed by the conditional probability $q(\mathbf{x}|\mathbf{y}; \boldsymbol{\zeta})$, its parameter is modified to fit well the corresponding Yang machine $p(\mathbf{y}|\mathbf{x}; \boldsymbol{\xi})$. As $\boldsymbol{\xi}$ changes by learning, $\boldsymbol{\zeta}$ changes correspondingly. The Ying machine can be used to improve the primitive representation \mathbf{x} by filling missing information and removing noise by using $\tilde{\mathbf{x}}$.

4.4 Bayesian Ying-Yang machine

Let us consider joint probability distributions $p(\mathbf{x}, \mathbf{y}) = q(\mathbf{x}, \mathbf{y})$ which are equal. Here, the Ying machine and the Yang machine are identical, and hence the machines are in perfect matching. Now consider \mathbf{y} as the parameters

to specify the distribution of \mathbf{x} , and we have a parameterized statistical model

$$M_{\text{Ying}} = \{p(\mathbf{x}|\mathbf{y})\}, \quad (167)$$

generated by the Ying mechanism. Here, the parameter \mathbf{y} is a random variable, and hence the Bayesian framework is introduced. The prior probability distribution of \mathbf{y} is

$$p(\mathbf{y}) = \int p(\mathbf{x}, \mathbf{y}) d\mathbf{x}. \quad (168)$$

The Bayesian framework to estimate \mathbf{y} from observed \mathbf{x} , or more rigorously, a number of independent observations, $\mathbf{x}_1, \dots, \mathbf{x}_n$, is as follows. Let $\bar{\mathbf{x}} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ be observed data and

$$p(\bar{\mathbf{x}}|\mathbf{y}) = \prod_{i=1}^n p(\mathbf{x}_i|\mathbf{y}). \quad (169)$$

The prior distribution of \mathbf{y} is $p(\mathbf{y})$, but after observing $\bar{\mathbf{x}}$, the posterior distribution is

$$p(\mathbf{y}|\bar{\mathbf{x}}) = \frac{p(\bar{\mathbf{x}}, \mathbf{y})}{\int p(\bar{\mathbf{x}}, \mathbf{y}) d\mathbf{y}}. \quad (170)$$

This is the Yang mechanism, and the

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} p(\mathbf{y}|\bar{\mathbf{x}}) \quad (171)$$

is the Bayesian posterior estimator.

A full exponential family $p(\mathbf{x}, \mathbf{y}) = \exp \{ \sum y_i x_i + k(\mathbf{x}) + w(\mathbf{y}) - \psi(\mathbf{x}, \mathbf{y}) \}$ is an interesting special case. Here, \mathbf{y} is the natural parameter for the exponential family of distributions,

$$q(\mathbf{x}|\mathbf{y}) = \exp \left\{ \sum y_i x_i + k(\mathbf{x}) - \psi(\mathbf{y}) \right\}. \quad (172)$$

The family of distributions $M_{\text{Ying}} = \{q(\mathbf{x}|\mathbf{y})\}$ specified by parameters \mathbf{y} is again an exponential family. It has a dually flat Riemannian structure.

Dually to the above, the family of the posterior distributions form a manifold $M_{\text{Yang}} = \{p(\mathbf{y}|\mathbf{x})\}$, consisting of

$$p(\mathbf{y}|\mathbf{x}) = \exp \left\{ \sum x_i y_i + w(\mathbf{y}) - \psi^*(\mathbf{y}) \right\}. \quad (173)$$

It is again an exponential family, where \mathbf{x} ($\bar{\mathbf{x}} = \sum \mathbf{x}_i/n$) is the natural parameter. It defines a dually flat Riemannian structure, too.

When the probability distribution $p(\mathbf{y})$ of the Yang machine includes a hyper parameter $p(\mathbf{y}) = p(\mathbf{y}, \boldsymbol{\zeta})$, we have a family $q(\mathbf{x}, \mathbf{y}; \boldsymbol{\zeta})$. It is possible that the Yang machine have a different parametric family $p(\mathbf{x}, \mathbf{y}; \boldsymbol{\xi})$. Then, we need to discuss the Ying mechanism and the Yang mechanism separately, keeping the Ying-Yang harmony.

This opens a new perspective to the Bayesian framework to be studied in future from the Ying-Yang standpoint. Information geometry gives a useful tool for analyzing it.

5 Geometry of belief propagation

A direct application of information geometry to machine learning is studied here. When a number of correlated random variables z_1, \dots, z_N exist, we need to treat their joint probability distribution $q(z_1, \dots, z_N)$. Here, we assume for simplicity's sake that z_i is binary, taking values 1 and -1 . Among all variables z_i , some are observed and the others are not. Therefore, z_i are divided into two parts, $\{x_1, \dots, x_n\}$ and $\{y_{n+1}, \dots, y_N\}$, where the values of $\mathbf{y} = (y_{n+1}, \dots, y_N)$ are observed. Our problem is to estimate the values of unobserved variables $\mathbf{x} = (x_1, \dots, x_n)$ based on observed \mathbf{y} . This is a typical problem of stochastic inference. We use the conditional probability distribution $q(\mathbf{x}|\mathbf{y})$ for estimating \mathbf{x} .

The graphical model or the Bayesian network is used to represent stochastic interactions among random variables [35]. In order to estimate \mathbf{x} , the belief propagation (BP) algorithm [15] uses a graphical model or a Bayesian network, and it provides a powerful algorithm in the field of artificial intelligence. However, its performance is not clear when the underlying graph has loopy structure. The present section studies this problem by using information geometry due to Ikeda, Tanaka and Amari [16,17]. The CCCP algorithm [36,37] is another powerful procedure for finding the BP solution. Its geometry is also studied.

5.1 Stochastic inference

We use a conditional probability $q(\mathbf{x}|\mathbf{y})$ to estimate the values of \mathbf{x} . Since \mathbf{y} is observed and fixed, we hereafter denote it simply by $q(\mathbf{x})$, suppressing observed variables \mathbf{y} , but it always means $q(\mathbf{x}|\mathbf{y})$, depending on \mathbf{y} .

We have two estimators of \mathbf{x} . One is the maximum likelihood estimator that maximizes $q(\mathbf{x})$:

$$\hat{\mathbf{x}}_{\text{mle}} = \arg \max_{\mathbf{x}} q(\mathbf{x}). \tag{174}$$

However, calculation of $\hat{\mathbf{x}}_{\text{mle}}$ is computationally difficult when n is large, because we need to search for the maximum among 2^n candidates \mathbf{x} 's. Another estimator $\tilde{\mathbf{x}}$ tries to minimize the expected number of errors of n component random variables $\tilde{x}_1, \dots, \tilde{x}_n$. When $q(\mathbf{x})$ is known, the expectation of x_i is

$$E[x_i] = \sum_{\mathbf{x}} x_i q(\mathbf{x}) = \sum x_i q_i(x_i). \tag{175}$$

This depends only on the marginal distribution of x_i ,

$$q_i(x_i) = \sum' q(x_1, \dots, x_n), \tag{176}$$

where summation \sum' is taken over all x_1, \dots, x_n except for x_i . The expectation is denoted by

$$\begin{aligned} \eta_i &= E[x_i] = \text{Prob}[x_i = 1] - \text{Prob}[x_i = -1] \\ &= q_i(1) - q_i(-1). \end{aligned} \tag{177}$$

When $\eta_i > 0$, $\text{Prob}[x_i = 1] > \text{Prob}[x_i = -1]$, so that, $\tilde{x}_i = 1$, otherwise $\tilde{x}_i = -1$. Therefore,

$$\tilde{x}_i = \text{sgn}(\eta_i) \tag{178}$$

is the estimator that minimizes the number of errors in the components of $\tilde{\mathbf{x}}$, where

$$\text{sgn}(u) = \begin{cases} 1, & u \geq 0, \\ -1, & u < 0. \end{cases} \tag{179}$$

The marginal distribution $q_i(x_i)$ is given in terms of η_i by

$$\text{Prob}\{x_i = 1\} = \frac{1 + \eta_i}{2}. \tag{180}$$

Therefore, our problem reduces to calculation of η_i , the expectation of x_i . However, exact calculation is computationally heavy when n is large, because (176) requires summation over 2^{n-1} \mathbf{x} 's. Physicists use the mean field approximation for this purpose [38]. The belief propagation is another powerful method of calculating it approximately by iterative procedures.

5.2 Graphical structure and random Markov structure

Let us consider a general probability distribution $p(\mathbf{x})$. In the special case where there are no interactions among x_1, \dots, x_n , they are independent, and the probability is written as the product of component distributions $p_i(x_i)$. Since we can always write the probabilities, in the binary case, as

$$p_i(1) = \frac{e^{h_i}}{\psi}, \quad p_i(-1) = \frac{e^{-h_i}}{\psi}, \tag{181}$$

where

$$\psi = e^{h_i} + e^{-h_i}, \tag{182}$$

we have

$$p_i(x_i) = \exp\{h_i x_i - \psi(h_i)\}. \tag{183}$$

Hence,

$$p(\mathbf{x}) = \exp\left\{\sum h_i x_i - \psi(\mathbf{h})\right\}, \tag{184}$$

where

$$\psi(\mathbf{h}) = \sum \psi(h_i), \tag{185}$$

$$\psi_i(h_i) = \log\{\exp(h_i) + \exp(-h_i)\}. \tag{186}$$

When x_1, \dots, x_n are not independent but their interactions exist, we denote the interaction among k variables x_{i_1}, \dots, x_{i_k} by a product term $x_{i_1} \cdot \dots \cdot x_{i_k}$,

$$c(\mathbf{x}) = x_{i_1} \cdot \dots \cdot x_{i_k}. \tag{187}$$

When there are many interaction terms, $p(\mathbf{x})$ is written as

$$p(\mathbf{x}) = \exp \left\{ \sum_i h_i x_i + \sum_r s_r c_r(\mathbf{x}) - \psi \right\}, \quad (188)$$

where ψ is the normalization factor, $c_r(\mathbf{x})$ represents a monomial interaction term such as

$$c_r(\mathbf{x}) = x_{i_1} \cdots x_{i_k}, \quad (189)$$

where r is an index showing a subset $\{i_1, \dots, i_k\}$ of indices $\{1, 2, \dots, n\}$, $k \geq 2$ and s_r is the intensity of such interaction.

When all interactions are pairwise, $k = 2$, we have $c_r(\mathbf{x}) = x_i x_j$ for $r = (i, j)$. To represent the structure of pairwise interactions, we use a graph $G = \{N, B\}$, in which we have n nodes N_1, \dots, N_n representing random variables x_1, \dots, x_n and b branches B_1, \dots, B_b . A branch B_r connects two nodes N_i and N_j , where r denotes the pair (x_i, x_j) of interaction. G is a non-directed graph having n nodes and b branches. When a graph has no loops, it is a tree graph, otherwise, it is a loopy graph. There are many physical and engineering problems having a graphical structure. They are, for example, spin glass, error-correcting codes, etc.

Interactions are not necessarily pairwise, and there may exist many interactions among more than two variables. Hence, we consider a general case (188), where r includes $\{i_1, \dots, i_k\}$, and b is the number of interactions. This is an example of the random Markov field, where each set $r = \{i_1, \dots, i_k\}$ forms a clique of the random graph.

We consider the following statistical model on graph G or a random Markov field specified by two vector parameters $\boldsymbol{\theta}$ and \mathbf{v} ,

$$M = \{p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{v})\}, \quad (190)$$

$$p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{v}) = \exp \left\{ \sum \theta_i x_i + \sum v_r c_r(\mathbf{x}) - \psi(\boldsymbol{\theta}, \mathbf{v}) \right\}. \quad (191)$$

This is an exponential family, forming a dually flat manifold, where $(\boldsymbol{\theta}, \mathbf{v})$ is its e -affine coordinates. Here, $\psi(\boldsymbol{\theta}, \mathbf{v})$ is a convex function from which a dually flat structure is derived. The model M includes the true distribution $q(\mathbf{x})$ in (13), i.e.,

$$q(\mathbf{x}) = \exp \{ \mathbf{h} \cdot \mathbf{x} + \mathbf{s} \cdot \mathbf{c}(\mathbf{x}) - \psi(\mathbf{h}, \mathbf{s}) \}, \quad (192)$$

which is given by $\boldsymbol{\theta} = \mathbf{h}$ and $\mathbf{v} = \mathbf{s} = (s_r)$. Our job is to find

$$\boldsymbol{\eta}^* = E_q[\mathbf{x}], \quad (193)$$

where expectation E_q is taken with respect to $q(\mathbf{x})$, and $\boldsymbol{\eta}$ is the dual (m -affine) coordinates corresponding to $\boldsymbol{\theta}$.

We consider $b+1$ e -flat submanifolds, M_0, M_1, \dots, M_b in M . M_0 is the manifold of independent distributions

given by $p_0(\mathbf{x}, \boldsymbol{\theta}) = \exp \{ \mathbf{h} \cdot \mathbf{x} + \boldsymbol{\theta} \cdot \mathbf{x} - \psi(\boldsymbol{\theta}) \}$, and its e -coordinates are $\boldsymbol{\theta}$.

For each branch or clique r , we also consider an e -flat submanifold M_r ,

$$M_r = \{p(\mathbf{x}, \boldsymbol{\zeta}_r)\}, \quad (194)$$

$$p(\mathbf{x}, \boldsymbol{\zeta}_r) = \exp \{ \mathbf{h} \cdot \mathbf{x} + s_r c_r(\mathbf{x}) + \boldsymbol{\zeta}_r \cdot \mathbf{x} - \psi_r(\boldsymbol{\zeta}_r) \}, \quad (195)$$

where $\boldsymbol{\zeta}_r$ is e -affine coordinates. Note that a distribution in M_r includes only one nonlinear interaction term $c_r(\mathbf{x})$. Therefore, it is computationally easy to calculate $E_r[\mathbf{x}]$ or $\text{Prob}\{x_i = 1\}$ with respect to $p(\mathbf{x}, \boldsymbol{\zeta}_r)$. All of M_0 and M_r play proper roles for finding $E_q[\mathbf{x}]$ of the target distribution (192).

5.3 m -projection

Let us first consider a special submanifold M_0 of independent distributions.

We show that the expectation $\boldsymbol{\eta}^* = E_q[\mathbf{x}]$ is given by the m -projection of $q(\mathbf{x}) \in M$ to M_0 [39,40]. We denote it by

$$\hat{p}(\mathbf{x}) = \prod_0 q(\mathbf{x}). \quad (196)$$

This is the point in M_0 that minimizes the KL-divergence from q to M_0 ,

$$\hat{p}(\mathbf{x}) = \arg \min_{p \in M_0} \text{KL}[q : p]. \quad (197)$$

It is known that \hat{p} is the point in M_0 such that the m -geodesic connecting q and \hat{p} is orthogonal to M_0 , and is unique, since M_0 is e -flat.

The m -projection has the following property.

Theorem 6 The m -projection of $q(\mathbf{x})$ to M_0 does not change the expectation of \mathbf{x} .

Proof Let $\hat{p}(\mathbf{x}) = p(\mathbf{x}, \boldsymbol{\theta}^*) \in M_0$ be the m -projection of $q(\mathbf{x})$. By differentiating $\text{KL}[q : p(\mathbf{x}, \boldsymbol{\theta})]$ with respect to $\boldsymbol{\theta}$, we have

$$\sum \mathbf{x} q(\mathbf{x}) - \frac{\partial}{\partial \boldsymbol{\theta}} \psi_0(\boldsymbol{\theta}^*) = 0. \quad (198)$$

The first term is the expectation of \mathbf{x} with respect to $q(\mathbf{x})$ and the second term is the $\boldsymbol{\eta}$ -coordinates of $\boldsymbol{\theta}^*$ which is the expectation of \mathbf{x} with respect to $p_0(\mathbf{x}, \boldsymbol{\theta}^*)$. Hence, they are equal.

5.4 Clique manifold M_r

The m -projection of $q(\mathbf{x})$ to M_0 is written as the product of the marginal distributions

$$\prod_0 q(\mathbf{x}) = \prod_{i=1}^n q_i(x_i), \quad (199)$$

which is what we want to obtain. Its coordinates are given by $\boldsymbol{\theta}^*$, from which we easily have $\boldsymbol{\eta}^* = E_{\boldsymbol{\theta}^*}[\mathbf{x}]$,

since M_0 consists of independent distributions. However, it is difficult to calculate the m -projection, or to obtain $\boldsymbol{\theta}^*$ or $\boldsymbol{\eta}^*$ directly. Physicists use the mean field approximation for this purpose. The belief propagation is another method of obtaining an approximation of $\boldsymbol{\theta}^*$. Here, the nodes and branches (cliques) play an important role.

Since the difficulty of calculations is given rise to by the existence of a large number of branches or cliques B_r , we consider a model M_r which includes only one branch or clique B_r . Since a branch (clique) manifold M_r includes only one nonlinear term $c_r(\mathbf{x})$, it is written as

$$p(\mathbf{x}, \boldsymbol{\zeta}_r) = \exp\{\mathbf{h} \cdot \mathbf{x} + s_r c_r(\mathbf{x}) + \boldsymbol{\zeta}_r \cdot \mathbf{x} - \psi_r\}, \quad (200)$$

where $\boldsymbol{\zeta}_r$ is a free vector parameter. Comparing this with $q(\mathbf{x})$, we see that the sum of all the nonlinear terms

$$\sum_{r' \neq r} s_{r'} c_{r'}(\mathbf{x}) \quad (201)$$

except for $s_r c_r(\mathbf{x})$ is replaced by a linear term $\boldsymbol{\zeta}_r \cdot \mathbf{x}$. Hence, by choosing an adequate $\boldsymbol{\zeta}_r$, $p(\mathbf{x}, \boldsymbol{\zeta}_r)$ is expected to give the same expectation $\mathbb{E}_q[\mathbf{x}]$ as $q(\mathbf{x})$ or its very good approximation. Moreover, M_r includes only one nonlinear term so that it is easy to calculate the expectation of \mathbf{x} with respect to $p(\mathbf{x}, \boldsymbol{\zeta}_r)$. In the following algorithm, all branch (clique) manifolds M_r cooperate iteratively to give a good entire solution.

5.5 Belief propagation

We have the true distribution $q(\mathbf{x})$, b clique distributions $p(\mathbf{x}, \boldsymbol{\zeta}_r) \in M_r$, $r = 1, \dots, b$, and an independent distribution $p_0(\mathbf{x}, \boldsymbol{\theta}) \in M_0$. All of them join force to search for a good approximation $\hat{p}(\mathbf{x}) \in M_0$ of $q(\mathbf{x})$.

Let $\boldsymbol{\zeta}_r$ be the current solution which M_r believes to give a good approximation of $q(\mathbf{x})$. We then project it to M_0 , giving the equivalent solution $\boldsymbol{\theta}_r$ in M_0 having the same expectation,

$$p(\mathbf{x}, \boldsymbol{\theta}_r) = \prod_0 p_r(\mathbf{x}, \boldsymbol{\zeta}_r). \quad (202)$$

We abbreviate this as

$$\boldsymbol{\theta}_r = \prod_0 \boldsymbol{\zeta}_r. \quad (203)$$

The $\boldsymbol{\theta}_r$ specifies the independent distribution equivalent to $p_r(\mathbf{x}, \boldsymbol{\zeta}_r)$ in the sense that they give the same expectation of \mathbf{x} . Since $\boldsymbol{\theta}_r$ includes both the effect of the single nonlinear term $s_r c_r(\mathbf{x})$ specific to M_r and that due to the linear term $\boldsymbol{\zeta}_r$ in (200),

$$\boldsymbol{\xi}_r = \boldsymbol{\theta}_r - \boldsymbol{\zeta}_r \quad (204)$$

represents the effect of the single nonlinear term $s_r c_r(\mathbf{x})$. This is the linearized version of $s_r c_r(\mathbf{x})$. Hence, M_r

knows that the linearized version of $s_r c_r(\mathbf{x})$ is $\boldsymbol{\xi}_r$, and broadcasts to all the other models M_0 and $M_{r'}$ ($r' \neq r$) that the linear counterpart of $s_r c_r(\mathbf{x})$ is $\boldsymbol{\xi}_r$. Receiving these messages from all M_r , M_0 guesses that the equivalent linear term of $\sum s_r c_r(\mathbf{x})$ will be

$$\boldsymbol{\theta} = \sum \boldsymbol{\xi}_r. \quad (205)$$

Since M_r in turn receives messages $\boldsymbol{\xi}_{r'}$ from all other $M_{r'}$, M_r uses them to form a new $\boldsymbol{\zeta}'_r$,

$$\boldsymbol{\zeta}'_r = \sum_{r' \neq r} \boldsymbol{\xi}_{r'}, \quad (206)$$

which is the linearized version of $\sum_{r' \neq r} s_{r'} c_{r'}(\mathbf{x})$. This process is repeated.

The above algorithm is written as follows, where the current candidates $\boldsymbol{\theta}^t \in M_0$ and $\boldsymbol{\zeta}_r^t \in M_r$ at time t , $t = 0, 1, \dots$, are renewed iteratively until convergence [17].

Geometrical BP Algorithm:

- 1) Put $t = 0$, and start with initial guesses $\boldsymbol{\zeta}_r^0$, for example, $\boldsymbol{\zeta}_r^0 = 0$.
- 2) For $t = 0, 1, 2, \dots$, m -project $p_r(\mathbf{x}, \boldsymbol{\zeta}_r^t)$ to M_0 , and obtain the linearized version of $s_r c_r(\mathbf{x})$,

$$\boldsymbol{\xi}_r^{t+1} = \prod_0 p_r(\mathbf{x}, \boldsymbol{\zeta}_r^t) - \boldsymbol{\zeta}_r^t. \quad (207)$$

- 3) Summarize all the effects of $s_r c_r(\mathbf{x})$, to give

$$\boldsymbol{\theta}^{t+1} = \sum_r \boldsymbol{\xi}_r^{t+1}. \quad (208)$$

- 4) Update $\boldsymbol{\zeta}_r$ by

$$\boldsymbol{\zeta}_r^{t+1} = \sum_{r' \neq r} \boldsymbol{\xi}_{r'}^{t+1} = \boldsymbol{\theta}^{t+1} - \boldsymbol{\xi}_r^{t+1}. \quad (209)$$

- 5) Stop when the algorithm converges.

5.6 Analysis of the solution of the algorithm

Let us assume that the algorithm has converged to $\{\boldsymbol{\zeta}_r^*\}$ and $\boldsymbol{\theta}^*$. Then, our solution is given by $p_0(\mathbf{x}, \boldsymbol{\theta}^*)$, from which

$$\eta_i^* = \mathbb{E}_0[x_i] \quad (210)$$

is easily calculated. However, this might not be the exact solution but is only an approximation. Therefore, we need to study its properties. To this end, we study the relations among the true distribution $q(\mathbf{x})$, the converged clique distributions $p_r^* = p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*)$ and the converged marginalized distribution $p_0^* = p_0(\mathbf{x}, \boldsymbol{\theta}^*)$. They are written as

$$q(\mathbf{x}) = \exp\left\{\mathbf{h} \cdot \mathbf{x} + \sum s_r c_r(\mathbf{x}) - \psi\right\}, \quad (211)$$

$$p_0(\mathbf{x}, \boldsymbol{\theta}^*) = \exp\left\{\mathbf{h} \cdot \mathbf{x} + \sum \boldsymbol{\xi}_r^* \cdot \mathbf{x} - \psi_0\right\}, \quad (212)$$

$$p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*) = \exp\left\{\mathbf{h} \cdot \mathbf{x} + \sum_{r' \neq r} \boldsymbol{\xi}_{r'}^* \cdot \mathbf{x} + s_r c_r(\mathbf{x}) - \psi_r\right\}. \quad (213)$$

The convergence point satisfies the two conditions:

$$1. \quad m\text{-condition} : \boldsymbol{\theta}^* = \prod_0 p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*); \quad (214)$$

$$2. \quad e\text{-condition} : \boldsymbol{\theta}^* = \frac{1}{b-1} \sum_{r=1}^b \boldsymbol{\zeta}_r^*. \quad (215)$$

Let us consider the m -flat submanifold M^* connecting $p_0(\mathbf{x}, \boldsymbol{\theta}^*)$ and all of $p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*)$,

$$M^* = \left\{ p(\mathbf{x}) \mid p(\mathbf{x}) = d_0 p_0(\mathbf{x}, \boldsymbol{\theta}^*) + \sum d_r p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*); \right. \\ \left. d_0 + \sum d_r = 1 \right\}. \quad (216)$$

This is a mixture family composed of $p_0(\mathbf{x}, \boldsymbol{\theta}^*)$ and $p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*)$.

We also consider the e -flat submanifold E^* connecting $p_0(\mathbf{x}, \boldsymbol{\theta}^*)$ and $p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*)$,

$$E^* = \left\{ p(\mathbf{x}) \mid \log p(\mathbf{x}) = v_0 \log p_0(\mathbf{x}, \boldsymbol{\theta}^*) \right. \\ \left. + \sum_r v_r \log p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*) - \psi \right\}. \quad (217)$$

Theorem 7 The m -condition implies that M^* intersects all of M_0 and M_r orthogonally. The e -condition implies that E^* include the true distribution $q(\mathbf{x})$.

Proof The m -condition guarantees that m -projection of p_r^* to M_0 is p_0^* . Since the m -geodesic connecting p_r^* and p_0 is included in M^* and the expectations are equal,

$$\boldsymbol{\eta}_r^* = \boldsymbol{\eta}_0^*, \quad (218)$$

and M^* is orthogonal to M_0 and M_r . The e -condition is easily shown by putting $c_0 = -(b-1)$, $c_r = 1$, because

$$q(\mathbf{x}) = c p_0(\mathbf{x}, \boldsymbol{\theta}^*)^{-(b-1)} \prod p_r(\mathbf{x}, \boldsymbol{\zeta}_r^*) \quad (219)$$

from (211)–(213), where $c = e^{-\psi_q}$.

The algorithm searches for $\boldsymbol{\theta}^*$ and $\boldsymbol{\zeta}_r^*$ until both the m - and e -conditions are satisfied. If M^* includes $q(\mathbf{x})$, its m -projection to M_0 is $p_0(\mathbf{x}, \boldsymbol{\theta}^*)$. Hence, $\boldsymbol{\theta}^*$ gives the true solution, but this is not guaranteed. Instead, E^* includes $q(\mathbf{x})$.

The following theorem is known and easy to prove.

Theorem 8 When the underlying graph is a tree, both E^* and M^* include $q(\mathbf{x})$, and the algorithm gives the exact solution.

5.7 CCCP procedure for belief propagation

The above algorithm is essentially the same as the BP algorithm given by Pearl [15] and studied by many followers. It is iterative search procedures for finding M^* and E^* . In steps 2)–5), it uses m -projection to obtain

new $\boldsymbol{\zeta}_r$'s, until the m -condition is satisfied eventually. The m -projections of all M_r become identical when the m -condition is satisfied. On the other hand, in each step of 3), we obtain $\boldsymbol{\theta} \in M_0$ such that the e -condition is satisfied. Hence, throughout the procedures, the e -condition is satisfied.

We have a different type of algorithm, which searches for $\boldsymbol{\theta}$ that eventually satisfies the e -condition, while the m -condition is always satisfied. Such an algorithm is proposed by Yuille [36] and Yuille and Rangarajan [37]. It consists of the two steps, beginning with initial guess $\boldsymbol{\theta}^0$:

CCCP Algorithm:

Step 1 (inner loop): Given $\boldsymbol{\theta}^t$, calculate $\{\boldsymbol{\zeta}_r^{t+1}\}$ by solving

$$\prod_0 p_r(\mathbf{x}, \boldsymbol{\zeta}_r^{t+1}) = b\boldsymbol{\theta}^t - \sum \boldsymbol{\zeta}_r^{t+1}. \quad (220)$$

Step 2 (outer loop): Given $\{\boldsymbol{\zeta}_r^{t+1}\}$, calculate $\boldsymbol{\theta}^{t+1}$ by

$$\boldsymbol{\theta}^{t+1} = b\boldsymbol{\theta}^t - \sum_r \boldsymbol{\zeta}_r^{t+1}. \quad (221)$$

The inner loop use an iterative procedures to obtain new $\{\boldsymbol{\zeta}_r^{t+1}\}$ in such a way that they satisfy the m -condition together with old $\boldsymbol{\theta}^t$. Hence, the m -condition is always satisfied, and we search for new $\boldsymbol{\theta}^{t+1}$ until the e -condition is satisfied.

We may simplify the inner loop by

$$\prod_0 p_r(\mathbf{x}, \boldsymbol{\zeta}_r^{t+1}) = p_0(\mathbf{x}, \boldsymbol{\theta}^t), \quad (222)$$

which can be solved directly. This gives a computationally easier procedure. There are some difference in the basin of attraction for the original and simplified procedures.

We have formulated a number of algorithms for stochastic reasoning in terms of information geometry. It not only clarifies the procedures intuitively but make it possible to analyze the stability of the equilibrium and speed of convergence. Moreover, we can estimate the error of estimation by using the curvatures of E^* and M^* . The new type of free energy is also defined. See Refs. [16,17] for more details.

6 Conclusions

We have given the dual geometrical structure in manifolds of probability distributions, positive measures or arrays, matrices, tensors and others. The dually flat structure is given from a convex function in general, which gives a Riemannian metric and a pair of dual flatness criteria. The information invariancy and dual flat structure are explained by using divergence functions without rigorous differential geometrical terminology. The dual geometrical structure is applied to various engineering problems, which include statistical inference, machine learning, optimization, signal processing

and neural networks. Applications to alternative optimization, Ying-Yang machine and belief propagations are shown from the geometrical point of view.

References

1. Amari S, Nagaoka H. *Methods of Information Geometry*. New York: Oxford University Press, 2000
2. Csiszár I. Information-type measures of difference of probability distributions and indirect observations. *Studia Scientiarum Mathematicarum Hungarica*, 1967, 2: 299–318
3. Bregman L. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 1967, 7(3): 200–217
4. Eguchi S. Second order efficiency of minimum contrast estimators in a curved exponential family. *The Annals of Statistics*, 1983, 11(3): 793–803
5. Chentsov N N. *Statistical Decision Rules and Optimal Inference*. Rhode Island, USA: American Mathematical Society, 1982 (originally published in Russian, Moscow: Nauka, 1972)
6. Dempster A P, Laird N M, Rubin D B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B*, 1977, 39(1): 1–38
7. Csiszár I, Tuszán G. Information geometry and alternating minimization procedures. *Statistics and Decisions*, 1984, Supplement Issue 1: 205–237
8. Amari S. Information geometry of the EM and em algorithms for neural networks. *Neural Networks*, 1995, 8(9): 1379–1408
9. Xu L. Bayesian Ying-Yang machine, clustering and number of clusters. *Pattern Recognition Letters*, 1997, 18(11–13): 1167–1178
10. Xu L. RBF nets, mixture experts, and Bayesian Ying-Yang learning. *Neurocomputing*, 1998, 19(1–3): 223–257
11. Xu L. Bayesian Kullback Ying-Yang dependence reduction theory. *Neurocomputing*, 1998, 22(1–3): 81–111
12. Xu L. BYY harmony learning, independent state space, and generalized APT financial analyses. *IEEE Transactions on Neural Networks*, 2001, 12(4): 822–849
13. Xu L. Best harmony, unified RPCL and automated model selection for unsupervised and supervised learning on Gaussian mixtures, three-layer nets and ME-RBF-SVM models. *International Journal of Neural Systems*, 2001, 11(1): 43–69
14. Xu L. BYY harmony learning, structural RPCL, and topological self-organizing on mixture models. *Neural Networks*, 2002, 15(8–9): 1125–1151
15. Pearl J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann, 1988
16. Ikeda S, Tanaka T, Amari S. Information geometry of turbo and low-density parity-check codes. *IEEE Transactions on Information Theory*, 2004, 50(6): 1097–1114
17. Ikeda S, Tanaka T, Amari S. Stochastic reasoning, free energy, and information geometry. *Neural Computation*, 2004, 16(9): 1779–1810
18. Csiszár I. Information measures: A critical survey. In: *Transactions of the 7th Prague Conference*. 1974, 83–86
19. Csiszár I. Axiomatic characterizations of information measures. *Entropy*, 2008, 10(3): 261–273
20. Ali M S, Silvey S D. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society. Series B*, 1966, 28(1): 131–142
21. Amari S. α -divergence is unique, belonging to both f -divergence and Bregman divergence classes. *IEEE Transactions on Information Theory*, 2009, 55(11): 4925–4931
22. Cichocki A, Adunek R, Phan A H, Amari S. *Nonnegative Matrix and Tensor Factorizations*. John Wiley, 2009
23. Havrda J, Charvát F. Quantification method of classification process: Concept of structural α -entropy. *Kybernetika*, 1967, 3: 30–35
24. Chernoff H. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 1952, 23(4): 493–507
25. Matsuyama Y. The α -EM algorithm: Surrogate likelihood maximization using α -logarithmic information measures. *IEEE Transactions on Information Theory*, 2002, 49(3): 672–706
26. Amari S. Integration of stochastic models by minimizing α -divergence. *Neural Computation*, 2007, 19(10): 2780–2796
27. Amari S. Information geometry and its applications: Convex function and dually flat manifold. In: Nielsen F ed. *Emerging Trends in Visual Computing. Lecture Notes in Computer Science, Vol 5416*. Berlin: Springer-Verlag, 2009, 75–102
28. Eguchi S, Copas J. A class of logistic-type discriminant functions. *Biometrika*, 2002, 89(1): 1–22
29. Murata N, Takenouchi T, Kanamori T, Eguchi S. Information geometry of U -boost and Bregman divergence. *Neural Computation*, 2004, 16(7): 1437–1481
30. Minami M, Eguchi S. Robust blind source separation by beta-divergence. *Neural Computation*, 2002, 14(8): 1859–1886
31. Byrne W. Alternating minimization and Boltzmann machine learning. *IEEE Transactions on Neural Networks*, 1992, 3(4): 612–620
32. Amari S, Kurata K, Nagaoka H. Information geometry of Boltzmann machines. *IEEE Transactions on Neural Networks*, 1992, 3(2): 260–271
33. Amari S. Natural gradient works efficiently in learning. *Neural Computation*, 1998, 10(2): 251–276
34. Amari S, Takeuchi A. Mathematical theory on formation of category detecting nerve cells. *Biological Cybernetics*, 1978, 29(3): 127–136
35. Jordan M I. *Learning in Graphical Models*. Cambridge, MA: MIT Press, 1999
36. Yuille A L. CCCP algorithms to minimize the Bethe and Kikuchi free energies: Convergent alternatives to belief propagation. *Neural Computation*, 2002, 14(7): 1691–1722
37. Yuille A L, Rangarajan A. The concave-convex procedure. *Neural Computation*, 2003, 15(4): 915–936
38. Oppen M, Saad D. *Advanced Mean Field Methods—Theory and Practice*. Cambridge, MA: MIT Press, 2001

39. Tanaka T. Information geometry of mean-field approximation. *Neural Computation*, 2000, 12(8): 1951–1968
40. Amari S, Ikeda S, Shimokawa H. Information geometry and mean field approximation: The α -projection approach. In: Oppen M, Saad D, eds. *Advanced Mean Field Methods—Theory and Practice*. Cambridge, MA: MIT Press, 2001, 241–257



Shun-ichi Amari was born in Tokyo, Japan, on January 3, 1936. He graduated from the Graduate School of the University of Tokyo in 1963 majoring in mathematical engineering and received Degree of Doctor of Engineering.

He worked as an Associate Professor at Kyushu University and the University of Tokyo, and then a Full Professor at the University of Tokyo, and is now Professor-Emeritus. He moved to RIKEN Brain Science Institute and served as Director for five years and is now Senior Advisor. He has been engaged in research in wide ar-

reas of mathematical science and engineering, such as topological network theory, differential geometry of continuum mechanics, pattern recognition, and information sciences. In particular, he has devoted himself to mathematical foundations of neural networks, including statistical neurodynamics, dynamical theory of neural fields, associative memory, self-organization, and general learning theory. Another main subject of his research is information geometry initiated by himself, which applies modern differential geometry to statistical inference, information theory, control theory, stochastic reasoning, and neural networks, providing a new powerful method to information sciences and probability theory.

Dr. Amari is past President of International Neural Networks Society and Institute of Electronic, Information and Communication Engineers, Japan. He received Emanuel R. Piore Award and Neural Networks Pioneer Award from the IEEE, the Japan Academy Award, C&C Award and Caianiello Memorial Award. He was the founding co-editor-in-chief of *Neural Networks*, among many other journals.