# The Centroids of Symmetrized Bregman Divergences 

Frank Nielsen ${ }^{1}$ Richard Nock ${ }^{2}$

${ }^{1}$ Sony Computer Science Laboratories Inc., FRL
École Polytechnique, LIX
Frank.Nielsen@acm.org
${ }^{2}$ University of Antilles-Guyanne
CEREGMIA
Richard.Nock@martinique.univ-ag.fr

## December 2007

## The centroid in Euclidean geometry

Given a point set $\mathcal{P}=\left\{p_{1}, \ldots, p_{n}\right\}$ of $\mathbb{E}^{d}$, the centroid $\bar{c}$ :

- Is the center of mass: $\bar{c}=\frac{1}{n} \sum_{i=1}^{n} p_{i}$,
- Minimizes $\min _{c \in \mathbb{R}^{d}} \sum_{i=1}^{n} \frac{1}{n}\left\|c p_{i}\right\|^{2}$ :

MinAvg squared Euclidean distance optimization,

- Plays a central role in center-based clustering methods (k-means of Lloyd'1957)


## Centroid and barycenters

Notion of centroid extends to barycenters:

$$
\bar{b}(w)=\sum_{i=1}^{n} w_{i} p_{i}
$$

$\rightarrow$ barycentric coordinates for interpolation, with $\|w\|=1$. Barycenter $\bar{b}(w)$ minimizes:

$$
\min _{c \in \mathbb{R}^{d}} \sum_{i=1}^{n} w_{i}\left\|c p_{i}\right\|^{2}
$$

$\rightarrow$ weighted MıNAVG optimization.
Intracluster min. weighted average $\sum_{i=1}^{n} w_{i}\left\|\bar{b}(w) p_{i}\right\|^{2}$.

## Center points in Euclidean geometry



- Centroid $\times$ : robust to outliers with simple closed-form sol. "Mean" radius is $\frac{1}{n} \sum_{i=1}^{n}\left\|\bar{c} p_{i}\right\|^{2}$.
- Circumcenter $\square$ : minimizes the radius of enclosing ball. Combinatorially defined by at most $d+1$ points. (MiniBall Welzl'1991)
MiniMAX (non-differentiable) optimization problem:

$$
C=\min _{c} \max _{i}\left\|c p_{i}\right\|^{2}
$$

- $\operatorname{MinAvG}\left(L_{2}\right) \rightarrow$ Fermat-Weber point $\circ$
$\rightarrow$ no closed form solution.


## Bregman divergences

Aim at generalizing Euclidean centroids to dually flat spaces.

Bregman divergences $D_{F}$ :
$F: \mathcal{X} \rightarrow \mathbb{R}$ strictly convex and differentiable function defined over an open convex domain $\mathcal{X}$ :
$D_{F}(p \| q)=F(p)-F(q)-\left\langle p-q, \nabla_{F}(q)\right\rangle$

- not a metric (symmetry and triangle inequality may fail)
- versatile family, popular in Comp. Sci.-Machine learning.


## Bregman divergence: Geometric interpretation



$$
\begin{array}{r}
D_{F}(p \| q)=F(p)-F(q)-\left\langle p-q, \nabla_{F}(q)\right\rangle \\
D_{F}(p \| q) \geq 0
\end{array}
$$

(with equality iff. $p=q$ )
$F$ : generator, contrast function or potential function (inf. geom.).

## Bregman divergence

## Example 1: The squared Euclidean distance

- $F(x)=x^{2}$ : strictly convex and differentiable over $\mathbb{R}^{d}$ (Multivariate $F(x)=\sum_{i=1}^{d} x_{i}^{2}$, obtained coordinatewise)

$$
\begin{aligned}
D_{F}(p \| q) & =F(p)-F(q)-\left\langle p-q, \nabla_{F}(q)\right\rangle \\
& =p^{2}-q^{2}-\langle p-q, 2 q\rangle \\
& -p^{2}-q^{2}-2\langle p, q\rangle+2 q^{2} \\
& =\|p-q\|^{2}
\end{aligned}
$$

## Example 1: The squared Euclidean distance



Java applet:

http://www.sonycsl.co.jp/person/nielsen/BregmanDivergence/

## Bregman divergence

## Example 2: The relative entropy (Kullback-Leibler divergence)

- $F(p)=\int p(x) \log p(x) \mathrm{d} x \quad$ (negative Shannon entropy) (Discrete distributions $F(p)=\sum_{x} p(x) \log p(x) \mathrm{d} x$ )

$$
\begin{aligned}
D_{F}(p \| q)= & \int(p(x) \log p(x)-q(x) \log q(x) \\
& -\langle p(x)-q(x), \log q(x)+1\rangle)) \mathrm{d} x \\
= & \int p(x) \log \frac{p(x)}{q(x)} \mathrm{d} x \quad(\text { KL divergence })
\end{aligned}
$$

Kullback-Leiber divergence also known as: relative entropy, discrimination information or $l$-divergence.
(Defined either on the probability simplex or extended on the full positive quadrant - unnormalized pdf.)

# Example 2: The relative entropy (Kullback-Leibler divergence) 



Java applet:
http://www.sonycsl.co.jp/person/nielsen/BregmanDivergence/

## Bregman divergences: A versatile family of measures

Bregman divergences are versatile, suited to mixed-type data.
(Build mixed-type multivariate divergences dimensionwise using elementary uni-type divergences.)

## Fact (Linearity)

Bregman divergence is a linear operator:
$\forall F_{1} \in \mathcal{C} \forall F_{2} \in \mathcal{C} \quad D_{F_{1}+\lambda F_{2}}(p \| q)=D_{F_{1}}(p \| q)+\lambda D_{F_{2}}(p \| q)$ for any $\lambda \geq 0$.

## Fact (Equivalence classes)

Let $G(x)=F(x)+\langle a, x\rangle+b$ be another strictly convex and differentiable function, with $a \in \mathbb{R}^{d}$ and $b \in \mathbb{R}$. Then $D_{F}(p \| q)=D_{G}(p \| q)$.
(Simplify Bregman generators by removing affine terms.)

## Sided and symmetrized centroids from MINAvg

Right-sided and left-sided centroids:

$$
\begin{aligned}
& c_{R}^{F}=\arg \min _{c \in \mathcal{X}} \frac{1}{n} \sum_{i=1}^{n} D_{F}\left(p_{i} \mid \| \mathrm{c}\right) \\
& c_{L}^{F}=\arg \min _{c \in \mathcal{X}} \frac{1}{n} \sum_{i=1}^{n} D_{F}\left(\boxed{c} \| \mid p_{i}\right)
\end{aligned}
$$

Symmetrized centroid:

$$
c^{F}=\arg \min _{c \in \mathcal{X}} \frac{1}{n} \sum_{i=1}^{n} \underbrace{\frac{D_{F}\left(p_{i} \mid \| \bar{c}\right)+D_{F}\left(\boxed{c} \| \mid p_{i}\right)}{2}}_{S_{F}}
$$

Symmetrized asymmetric Bregman divergence $S_{F}$ is not a Bregman divergence (because of domain convexity)

## Sided right-type centroid

## Theorem

The right-type sided Bregman centroid $c_{R}^{F}$ of a set $\mathcal{P}$ of $n$ points $p_{1}, \ldots p_{n}$, defined as the minimizer for the average right divergence
$c_{R}^{F}=\arg \min _{c} \sum_{i=1}^{n} \frac{1}{n} D_{F}\left(p_{i} \| c\right)=\arg \min _{c} \operatorname{AVG}_{F}(\mathcal{P} \| c)$, is unique, independent of the selected divergence $D_{F}$, and coincides with the center of mass $c_{R}^{F}=c_{R}=\bar{p}=\frac{1}{n} \sum_{i=1}^{n} p_{i}$.

## Proof:

Start with the function to minimize:

$$
\operatorname{AVG}_{F}(\mathcal{P} \| q)=\sum_{i=1}^{n} \frac{1}{n} D_{F}\left(p_{i} \| q\right)
$$

$$
\begin{aligned}
\operatorname{AVG}_{F}(\mathcal{P} \| q) & =\sum_{i=1}^{n} \frac{1}{n}\left(F\left(p_{i}\right)-F(q)-<p_{i}-q, \nabla F(q)>\right) \\
& \begin{aligned}
\operatorname{AVG}_{F}(\mathcal{P}, q) & =\left(\sum_{i=1}^{n} \frac{1}{n} F\left(p_{i}\right)-F(\bar{p})\right)+\left(F(\bar{p})-F(q)-\sum_{i=1}^{n} \frac{1}{n}\left\langle p_{i}-q, \nabla F(q)\right\rangle\right), \\
& \left(\sum_{i=1}^{n} \frac{1}{n} F\left(p_{i}\right)-F(\bar{p})\right)+\left(F(\bar{p})-F(q)-\left\langle\sum_{i=1}^{n} \frac{1}{n}\left(p_{i}-q\right), \nabla F(q)\right\rangle\right), \\
& =\underbrace{\left(\frac{1}{n} \sum_{i=1}^{n} F\left(p_{i}\right)-F(\bar{p})\right)}_{\text {Independent of } q}+D_{F}(\bar{p}| | q) .
\end{aligned}
\end{aligned}
$$

- Minimized for $q=\bar{p}$ since $D_{F}(\bar{p} \| q) \geq 0$ with equality iff $q=\bar{p}$
- Information radius: $\mathrm{JS}_{F}(\mathcal{P})=\frac{1}{n} \sum_{i=1}^{n} F\left(p_{i}\right)-F(\bar{p}) \geq 0$ $\rightarrow$ Jensen $F$ remainder aka. Burbea-Rao divergences $\leftarrow$ Generalize Jensen-Shannon divergences (Lin'1991)


## Dual Bregman divergence

Legendre-Fenchel (slope) transformation $F \rightarrow G=\mathcal{L} F$ :

$$
G(y)=\sup _{x \in \mathcal{X}}\{<y, x>-F(x)\}
$$

$G=F^{*}\left(\right.$ with $\left.G^{*}=F^{* *}=F\right)$, minimized for $y=x^{\prime} \stackrel{\text { def }}{=} \nabla F(x)$.

$$
F^{*}\left(x^{\prime}\right)=<x, x^{\prime}>-F(x)
$$

## Dual Bregman divergence

$$
\begin{aligned}
D_{F}(p \| q) & =F(p)+F^{*}(\nabla F(q))-<p, \nabla F(q)> \\
& =F(p)+F^{*}\left(q^{\prime}\right)-<p, q^{\prime}> \\
& =D_{F^{*}}\left(q^{\prime} \| p^{\prime}\right)
\end{aligned}
$$

More details in "Breaman Voronoi diaarams", arXiv:0709.2196

## Left-sided Bregman centroid

## Theorem

The left-type sided Bregman centroid $c_{L}^{F}$, defined as the minimizer for the average left divergence $c_{L}^{F}=\arg \min _{c \in \mathcal{X}} \operatorname{AVG}_{L}^{F}(c \| \mathcal{P})$, is the unique point $c_{L}^{F} \in \mathcal{X}$ such that $c_{L}^{F}=(\nabla F)^{-1}\left(\overline{p^{\prime}}\right)=(\nabla F)^{-1}\left(\sum_{i=1}^{n} \nabla F\left(p_{i}\right)\right)$, where $\overline{p^{\prime}}=c_{R}^{F^{*}}\left(\mathcal{P}_{F^{\prime}}\right)$ is the center of mass for the gradient point set $\mathcal{P}_{F}^{\prime}=\left\{p_{i}^{\prime}=\nabla F\left(p_{i}\right) \mid p_{i} \in \mathcal{P}\right\}$.
$c_{L}^{F}=\arg \min _{c \in \mathcal{X}} \operatorname{AVG}_{F}(\boxed{\mathrm{c}} \| \mathcal{P}) \Leftrightarrow \arg \min _{C^{\prime} \in \mathcal{X}}^{\prime} \operatorname{AVG}_{F^{*}}\left(\mathcal{P}_{F^{\prime}}^{\prime} \| \boxed{\mathrm{c}^{\prime}}\right)=c_{R}^{F^{*}}\left(\mathcal{P}_{F}^{\prime}\right)$
The information radii of sided Bregman centroids are equal:
$\operatorname{AVG}_{F}\left(\mathcal{P} \| c_{R}^{F}\right)=\operatorname{AVG}_{F}\left(c_{L}^{F} \| \mathcal{P}\right)=\operatorname{JS}_{F}(\mathcal{P})=\frac{1}{n} \sum_{i=1}^{n} F\left(p_{i}\right)-F(\bar{p})>0$
is the $F$-Jensen-Shannon divergence for the uniform weight $\bar{\equiv}$

## Generalized means

A sequence $\mathcal{V}$ of $n$ real numbers $V=\left\{v_{1}, \ldots, v_{n}\right\}$

$$
M(\mathcal{V} ; f)=f^{-1}\left(\frac{1}{n} \sum_{i=1}^{n} f\left(v_{i}\right)\right)
$$

For example, Pythagoras' means:

- Arithmetic: $f(x)=x$ (average)
- Geometric: $f(x)=\log x$ (central tendency)
- Harmonic: $f(x)=\frac{1}{x}$ (average of rates)

$$
\min _{i} x_{i} \leq M(\mathcal{V} ; f) \leq \max _{i} x_{i}
$$

min and max: power means $\left(f(x)=x^{p}\right)$ for $p \rightarrow \pm \infty$

## Bijection: Bregman divergences and means

Bijection: Bregman divergence $D_{F} \leftrightarrow \nabla F$-means

$$
M(\mathcal{S} ; f)=M(\mathcal{S} ; a f+b) \forall a \in \mathbb{R}_{*}^{+} \text {and } \forall b \in \mathbb{R}
$$

Recall one property of Bregman divergences:

## Fact (Equivalence classes)

Let $G(\mathbf{x})=F(\mathbf{x})+\langle\mathbf{a}, \mathbf{x}\rangle+b$ be another strictly convex and differentiable function, with $\mathbf{a} \in \mathbb{R}^{d}$ and $b \in \mathbb{R}$. Then $D_{F}(\mathbf{p} \| \mathbf{q})=D_{G}(\mathbf{p} \| \mathbf{q})$.

## Left- and right-sided Bregman barycenters

Left- and right-sided Bregman centroids extend to barycenters.

## Theorem

Bregman divergences are in bijection with generalized means. The right-type barycenter $b_{R}^{F}(w)$ is independent of $F$ and computed as the weighted arithmetic mean on the point set, a generalized mean for the identity function:
$b_{R}^{F}(\mathcal{P} ; w)=b_{R}(\mathcal{P} ; w)=M(\mathcal{P} ; x ; w)$ with
$M(\mathcal{P} ; f ; w) \stackrel{\text { def }}{=} f^{-1}\left(\sum_{i=1}^{n} w_{i} f\left(v_{i}\right)\right)$. The left-type Bregman
barycenter $b_{L}^{F}$ is computed as a generalized mean on the point set for the gradient function: $b_{L}^{F}(\mathcal{P})=M(\mathcal{P} ; \nabla F ; w)$.
The information radius of sided barycenters is:
$\operatorname{JS}_{F}(\mathcal{P} ; w)=\sum_{i=1}^{d} w_{i} F\left(p_{i}\right)-F\left(\sum_{i=1}^{d} w_{i} p_{i}\right)$.

## Symmetrized Bregman centroid

Symmetrized centroid defined from the minimum average optimization problem:

$$
c^{F}=\arg \min _{c \in \mathcal{X}} \frac{1}{n} \sum_{i=1}^{n} \frac{D_{F}\left(c \| p_{i}\right)+D_{F}\left(p_{i} \| c\right)}{2}=\arg \min _{c \in \mathcal{X}} \operatorname{AVG}(\mathcal{P} ; c)
$$

## Lemma

The symmetrized Bregman centroid $c^{F}$ is unique and obtained by minimizing $\min _{q \in \mathcal{X}} D_{F}\left(c_{R}^{F} \| q\right)+D_{F}\left(q \| c_{L}^{F}\right)$ :

$$
c^{F}=\arg \min _{q \in \mathcal{X}} D_{F}\left(c_{R}^{F} \| q\right)+D_{F}\left(q \| c_{L}^{F}\right) .
$$

$\rightarrow$ Minimization problem depends only on sided centroids.

$$
\begin{aligned}
\operatorname{AVG}_{F}(\mathcal{P} \| q) & =\left(\sum_{i=1}^{n} \frac{1}{n} F\left(p_{i}\right)-F(\bar{p})\right)+D_{F}(\bar{p} \| q) \\
\operatorname{AVG}_{F}(q \| \mathcal{P}) & =\operatorname{AVG}_{F^{*}}\left(\mathcal{P}_{F^{\prime}} \| q^{\prime}\right) \\
& =\left(\sum_{i=1}^{n} \frac{1}{n} F^{*}\left(p_{i}^{\prime}\right)-F^{*}\left(\overline{p^{\prime}}\right)\right)+D_{F^{*}}\left(\overline{p_{F}^{\prime}} \| q_{F}^{\prime}\right)
\end{aligned}
$$

But $D_{F^{*}}\left(\overline{p_{F}^{\prime}} \| q_{F}^{\prime}\right)=D_{F^{* *}}\left(\nabla F^{*} \circ \nabla F(q) \| \nabla F^{*}\left(\sum_{i=1}^{n} \nabla F\left(p_{i}\right)\right)\right)=$ $D_{F}\left(q \| c_{L}^{F}\right)$ since $F^{* *}=F, \nabla F^{*}=\nabla F^{-1}$ and $\nabla F^{*} \circ \nabla F(q)=q$.
$\arg \min _{c \in \mathcal{X}} \frac{1}{2}\left(\operatorname{AVG}_{F}(\mathcal{P} \| q)+\operatorname{AVG}_{F}(q \| \mathcal{P})\right) \Longleftrightarrow \arg \min _{q \in \mathcal{X}} D_{F}\left(c_{R}^{F} \| q\right)+D_{F}\left(q \| c_{L}^{F}\right)$
(removing all terms independent of $q$ )

## Geometric characterization

## Theorem

The symmetrized Bregman centroid $c^{F}$ is uniquely defined as the minimizer of $D_{F}\left(c_{R}^{F} \| q\right)+D_{F}\left(q \| c_{L}^{F}\right)$. It is defined geometrically as $c^{F}=\Gamma_{F}\left(c_{R}^{F}, c_{L}^{F}\right) \cap M_{F}\left(c_{R}^{F}, c_{L}^{F}\right)$, where $\Gamma_{F}\left(c_{R}^{F}, c_{L}^{F}\right)=\left\{(\nabla F)^{-1}\left((1-\lambda) \nabla F\left(c_{R}^{F}\right)+\lambda \nabla F\left(c_{L}^{F}\right)\right) \mid \lambda \in[0,1]\right\}$ is the geodesic linking $c_{R}^{F}$ to $c_{L}^{F}$, and $M_{F}\left(c_{R}^{F}, c_{L}^{F}\right)$ is the mixed-type Bregman bisector:

$$
M_{F}\left(c_{R}^{F}, c_{L}^{F}\right)=\left\{x \in \mathcal{X} \mid D_{F}\left(c_{R}^{F} \| x\right)=D_{F}\left(x \| c_{L}^{F}\right)\right\} .
$$

Proof by contradiction using Bregman Pythagoras' theorem.

## Generalized Bregman Pythagoras' theorem



Projection $\mathbf{p}_{\mathcal{W}}$ of point p to a convex subset $\mathcal{W} \subseteq \mathcal{X}$.

$$
D_{F}(\mathbf{w} \| \mathbf{p}) \geq D_{F}\left(\mathbf{w} \| \mathbf{p}_{\mathcal{W}}\right)+D_{F}\left(\mathbf{p}_{\mathcal{W}} \| \mathbf{p}\right)
$$

with equality for and only for affine sets $\mathcal{W}$


Bregman projection: $q_{\perp}=\arg \min _{t \in \Gamma\left(c_{R}^{F}, c_{L}^{\mathcal{C}_{L}^{\prime}}\right)} D_{F}(t \| q)$

$$
\begin{aligned}
D_{F}(p \| q) & \geq D\left(p \| P_{\Omega}(q)\right)+D_{F}\left(P_{\Omega}(q) \| q\right) \\
P_{\Omega}(q) & =\arg \min _{\omega \in \Omega} D_{F}(\omega \| q) \\
D_{F}\left(c_{R}^{F} \| q\right) & \geq D_{F}\left(c_{R} \| q_{\perp}\right)+D_{F}\left(q_{\perp} \| q\right) \\
D_{F}\left(q \| c_{L}^{F}\right) & \geq D_{F}\left(q \| q_{\perp}\right)+D_{F}\left(q_{\perp} \| C_{L}^{F}\right)
\end{aligned}
$$

$$
D_{F}\left(c_{R}^{F} \| q\right)+D_{F}\left(q \| c_{L}^{F}\right) \geq D_{F}\left(c_{R}^{F} \| q_{\perp}\right)+D_{F}\left(q_{\perp} \| c_{L}^{F}\right)+\left(D_{F}\left(q_{\perp} \| q\right)+D_{F}\left(q \| q_{\perp}\right)\right)
$$

But $D_{F}\left(q_{\perp} \| q\right)+D_{F}\left(q \| q_{\perp}\right)>0$ yields contradiction.
Proved mixed-type bisector by moving $q$ along the geodesic while reducing $\left|D_{F}\left(c_{R}^{F} \| q\right)-D_{F}\left(q \| c_{L}^{F}\right)\right|$.

## A few examples: Kullback-Leibler \& Itakura-Saito



## A few examples: Exponential \& Logistic losses



## Non-linear mixed-type bisector

$$
\begin{aligned}
M_{F}(p, q)=\{ & x \in \mathcal{X} \mid F(p)-F(q)-2 F(x)- \\
& \left.<p, x^{\prime}>+<x, x^{\prime}>+<x, q^{\prime}>-<q, q^{\prime}>=0\right\}
\end{aligned}
$$

$\rightarrow$ non-closed form solution.

## Corollary

The symmetrized Bregman divergence minimization problem is both lower and upper bounded as follows:

$$
\operatorname{JS}_{F}(\mathcal{P}) \leq \operatorname{AVG}_{F}\left(\mathcal{P} ; c^{F}\right) \leq D_{F}\left(c_{R}^{F} \| c_{L}^{F}\right)
$$

## Geodesic-walk dichotomic approximation algorithm

## ALGORITHM:

Geodesic is parameterized by $\lambda \in[0,1]$.
Start with $\lambda_{m}=0$ and $\lambda_{M}=1$.
Geodesic walk. Compute interval midpoint $\lambda_{h}=\frac{\lambda_{m}+\lambda_{M}}{2}$ and corresponding geodesic point

$$
q_{\lambda_{h}}=(\nabla F)^{-1}\left(\left(1-\lambda_{h}\right) \nabla F\left(c_{R}^{F}\right)+\lambda_{h} \nabla F\left(c_{L}^{F}\right)\right),
$$

Mixed-type bisector side. Evaluate the sign of

$$
D_{F}\left(c_{R}^{F} \| q_{h}\right)-D_{F}\left(q_{h} \| c_{L}^{R}\right), \text { and }
$$

Dichotomy. Branch on $\left[\lambda_{h}, \lambda_{M}\right]$ if the sign is negative, or on [ $\lambda_{m}, \lambda_{h}$ ] otherwise.

Precision and number of iterations as a function of $h_{F}=\max _{x \in \Gamma\left(c_{R}^{F}, c_{L}^{F}\right)}\left\|H_{F}(x)\right\|^{2}$. (Bregman balls, ECML'05)

## Applications

Applications of the dichotomic geodesic walk algorithm

- Symmetrized non-parametric Kullback-Leibler (SKL),
- Symmetrized Kullback-Leibler of multivariate normals, (parametric distributions)
- Symmetrized Bregman-Csiszár centroids, $\rightarrow$ include J-divergence and COSH distance (Itakura-Saito symmetrized divergence).


## The symmetrized Kullback-Leibler divergence

For two discrete probability mass functions $p$ and $q$ :

$$
K L(p \| q)=\sum_{i=1}^{d} p^{(i)} \log \frac{p^{(i)}}{q^{(i)}}
$$

For continuous distributions

$$
\operatorname{KL}(p \| q)=\int_{x} p(x) \log \frac{p(x)}{q(x)} .
$$

Symmetric Kullback-Leibler is called $J$-divergence.
Observe that finite discrete distributions are parametric.
$\rightarrow$ Degree of freedom is cardinal of sample space minus one.

## Exponential families and divergences

Exponential families in statistics have the following pdf.:

$$
\exp (<\theta, t(x)>-F(\theta)+C(x))
$$

- $\theta$ : natural parameter (dimension: order of the exp. fam.)
- $t(x)$ : sufficient statistics ( $\leftarrow$ Fisher-Neyman factorization),
- F: log normalizer (cumulant) characterizes the family,
- $C(x)$ : carrier measure (usually Lebesgue or counting).

Kullback-Leibler of exp. fam. is a Bregman divergence

$$
\operatorname{KL}\left(p\left(\theta_{p} \mid F\right) \| q\left(\theta_{q} \mid F\right)\right)=D_{F}\left(\theta_{q} \| \theta_{p}\right)
$$

## Discrete distributions are multinomials

Multinomials extend Bernoulli distributions.
Conversion from source $q$ to natural $\theta$ parameters:

$$
\theta^{(i)}=\log \frac{q^{(i)}}{q^{(d)}}=\log \frac{q^{(i)}}{1-\sum_{j=1}^{d-1} q^{(j)}}
$$

with $q^{(d)}=1-\sum_{j=1}^{d-1} q^{(j)}$
... and to convert back from natural to source parameters:

$$
q^{(i)}=\frac{\exp \theta^{(i)}}{1+\sum_{j=1}^{d-1}\left(1+\exp \theta^{(j)}\right)}
$$

with $q^{(d)}=\frac{1}{1+\sum_{j=1}^{d-1}\left(1+\exp \theta^{(j)}\right)}$

## Dual log normalizers

Log normalizer of multinomials is

$$
F(\theta)=\log \left(1+\sum_{i=1}^{d-1} \exp \theta^{(i)}\right)
$$

Logistic entropy on open space $\Theta=\mathbb{R}^{d-1}$.

Dual Legendre function $F^{*}=\mathcal{L} F$ is $d$-ary entropy:

$$
F^{*}(\eta)=\left(\sum_{i=1}^{d-1} \eta^{(i)} \log \eta^{(i)}\right)+\left(1-\sum_{i=1}^{d-1} \eta^{(i)}\right) \log \left(1-\sum_{i=1}^{d-1} \eta^{(i)}\right)
$$

## Geodesic of multinomials

Both $\nabla F$ and $\nabla F^{*}$ are necessary for:

- computing the left-sided centroid ( $\nabla F$-means),
- walking on the geodesic $\left((\nabla F)^{-1}=\nabla F^{*}\right)$ :

$$
q_{\lambda}=(\nabla F)^{-1}\left((1-\lambda) \nabla F\left(c_{R}^{F}\right)+\lambda \nabla F\left(c_{L}^{F}\right)\right)
$$

$$
\nabla F(\theta)=\left(\frac{\exp \theta^{(i)}}{1+\sum_{j=1}^{d-1} \exp \theta^{(j)}}\right)_{i}^{\stackrel{\text { def }}{=}} \eta
$$

It follows from Legendre transformation that $\nabla F^{*}=\nabla F^{-1}$

$$
\nabla^{-1} F(\eta)=\left(\log \frac{\eta^{(i)}}{1-\sum_{j=1}^{d-1} \eta^{(j)}}\right)_{i}^{\stackrel{\operatorname{def}}{=}} \theta
$$

## Example: Centroids of histograms

## Centroids are generalized means:

$\rightarrow$ coordinates are always inside the extrema...


## Example: Centroids of histograms

... but not necessarily true for the corresponding histograms!



## Centroids of histograms: Java Applet

Generalize ad-hoc convex programming method of (Veldhuis'02).




Try with your own images online using Java applet at:
http://www.sonycsl.co.jp/person/nielsen/§BDj/号

## Side-by-side comparison

## INPUT:

$n$ discrete distributions $q_{1}, \ldots, q_{n}$ of $\mathcal{S}^{d}$ with $\forall i \in\{1, \ldots, n\} \quad q_{i}=\left(q_{i}^{(1)}, \ldots, q_{i}^{(d)}\right)$.

## INITIALIZATION

Arithmetic mean:
$\forall k \bar{q}^{(k)}=\frac{1}{n} \sum_{i=1}^{n} q_{i}^{(k)}$
Geometric normalized mean:
$\forall k \breve{q}^{(k)}=\frac{\tilde{q}^{(k)}}{\sum_{i=1}^{d} \tilde{q}_{i}}$ with $\forall k \tilde{q}^{(k)}=\left(\prod_{i=1}^{n} q_{i}^{(k)}\right)^{\frac{1}{n}}$
$\alpha=-1$

## MAIN LOOP:

For 1 to 10

$$
\forall k y^{(k)}=\frac{\bar{q}^{(k)}}{\bar{q}^{(k)} \exp \alpha}
$$

$$
\forall k x^{(k)}=1
$$

INNER LOOP 1:
For 1 to 5

$$
\forall k x^{(k)} \leftarrow x^{(k)}-\frac{x^{(k)} \log x^{(k)}-y^{(k)}}{\log x^{(k)}+1}
$$

INNER LOOP 2:
For 1 to 5

$$
\alpha \leftarrow \alpha-\frac{\left(\sum_{i=1}^{d} x^{(k)} \breve{q}^{(k)} \exp \alpha\right)-1}{\sum_{i=1}^{d} x^{(k)} \breve{q}^{(k)} \exp \alpha}
$$

CENTROID:
$\forall k c^{(k)}=x^{(k)} \breve{q}^{(k)} \exp \alpha$
(Vedlhuis'02)

## INPUT:

$n$ discrete distributions $q_{1}, \ldots, q_{n}$ of $\mathcal{S}^{d}$ with
$\forall i \in\{1, \ldots, n\} q_{i}=\left(q_{i}^{(1)}, \ldots, q_{i}^{(d)}\right)$
CONVERSION:
Probability mass function $\rightarrow$ multinomial
$\forall i \forall k \theta_{i}^{(k)}=\log \frac{q_{i}^{(k)}}{1-\sum_{i=1}^{d-1} q_{i}^{(j)}}$
$F(\theta)=\log \left(1+\sum_{j=1}^{d-1} \exp \theta^{(j)}\right)$
$\nabla F(\theta)=\left(\frac{\exp ^{(i)}}{1+\sum_{j=1}^{d-1} \exp \theta^{(j)}}\right)_{i}$
$(\nabla F)^{-1(\eta)}=\left(\log \frac{\eta^{(i)}}{1-\sum_{i=1}^{d-1} \eta^{(j)}}\right)_{i \in\{1, \ldots, d-1\}}$

## INITIALIZATION:

Arithmetic mean: $\theta_{R}^{F}=\frac{1}{n} \sum_{i=1}^{n} \theta_{i}$
$\nabla F$-mean: $\theta_{L}^{F}=\nabla F^{-1}\left(\frac{1}{n} \sum_{i=1}^{n} \nabla F\left(\theta_{i}\right)\right)$
$\lambda_{m}=0, \lambda_{M}=1$
GEODESIC DICHOTOMIC WALK:
While $\lambda_{M}-\lambda_{m}>$ precision do

$$
\begin{aligned}
& \lambda=\frac{\lambda_{m}+\lambda_{M}}{2} \\
& \theta=(\nabla F)^{-1}\left((1-\lambda) \nabla F\left(c_{R}^{F}\right)+\lambda \nabla F\left(c_{L}^{F}\right)\right) \\
& \text { if } D_{F}\left(c_{R}^{F} \| \theta\right)>D_{F}\left(\theta \| c_{L}^{F}\right) \text { then } \\
& \\
& \qquad \begin{array}{l}
\lambda_{M}=\lambda \\
\\
\\
\\
\\
\\
\end{array} \lambda_{m}=\lambda
\end{aligned}
$$

CONVERSION:
Multinomial $\rightarrow$ Probability mass function
$\forall i q_{i}^{(d)}=\frac{1}{1+\sum_{j=1}^{d-1}\left(1+\exp \theta_{i}^{(j)}\right)}$
$\forall i \forall k q_{i}^{(k)}=\frac{\exp \theta_{i}^{(k)}}{1+\sum_{j=1}^{d-1}\left(1+\exp \theta_{i}^{(j)}\right)}$
Geodesic dichotomic walk

## Entropic means of multivariate normal distributions

Multivariate normal distributions of $\mathbb{R}^{d}$ has following pdf.:

$$
\begin{aligned}
\operatorname{Pr}(X=x) & =p(x ; \mu, \Sigma) \\
& =\frac{1}{(2 \pi)^{\frac{d}{2}} \sqrt{\operatorname{det} \Sigma}} \exp \left(-\frac{(x-\mu)^{\top} \Sigma^{-1}(x-\mu)}{2}\right)
\end{aligned}
$$

Source parameter is a mixed-type of vector and matrix:

$$
\tilde{\Lambda}=(\mu, \Sigma)
$$

Order of the parametric distribution family is $D=\frac{d(d+3)}{2}>d$.

## Exponential family decomposition

Multivariate normal distribution belongs to the exponential families:

$$
\exp (<\theta, t(x)>-F(\theta)+C(x))
$$

- Sufficient statistics: $\tilde{x}=\left(x,-\frac{1}{2} x x^{\top}\right)$
- Natural parameters: $\tilde{\Theta}=(\theta, \Theta)=\left(\Sigma^{-1} \mu, \frac{1}{2} \Sigma^{-1}\right)$
- Log normalizer

$$
F(\tilde{\Theta})=\frac{1}{4} \operatorname{Tr}\left(\Theta^{-1} \theta \theta^{T}\right)-\frac{1}{2} \log \operatorname{det} \Theta+\frac{d}{2} \log \pi
$$

Mixed-type separable inner product:

$$
<\tilde{\Theta}_{p}, \tilde{\Theta}_{q}>=<\Theta_{p}, \Theta_{q}>+<\theta_{p}, \theta_{q}>
$$

with matrix inner product defined as the following trace:

$$
<\Theta_{p}, \Theta_{q}>=\operatorname{Tr}\left(\Theta_{p} \Theta_{q}^{T}\right)
$$

## Dual Legendre functions

$$
\begin{aligned}
F(\tilde{\Theta}) & =\frac{1}{4} \operatorname{Tr}\left(\Theta^{-1} \theta \theta^{T}\right)-\frac{1}{2} \log \operatorname{det} \Theta+\frac{d}{2} \log \pi \\
F^{*}(\tilde{H}) & =-\frac{1}{2} \log \left(1+\eta^{T} H^{-1} \eta\right)-\frac{1}{2} \log \operatorname{det}(-H)-\frac{d}{2} \log (2 \pi e)
\end{aligned}
$$

Parameters transformations $\tilde{H} \leftrightarrow \tilde{\Theta} \leftrightarrow \Lambda$

$$
\begin{gathered}
\tilde{H}=\binom{\eta=\mu}{H=-\left(\Sigma+\mu \mu^{T}\right)} \Longleftrightarrow \tilde{\Lambda}=\binom{\lambda=\mu}{\Lambda=\Sigma} \Longleftrightarrow \tilde{\Theta}=\binom{\theta=\Sigma^{-1} \mu}{\Theta=\frac{1}{2} \Sigma^{-1}} \\
\tilde{H}=\nabla_{\tilde{\Theta}} F(\tilde{\Theta})=\binom{\nabla_{\tilde{\Theta}} F(\theta)}{\nabla_{\tilde{\Theta}} F(\Theta)}=\binom{\frac{1}{2} \Theta^{-1} \theta}{-\frac{1}{2} \Theta^{-1}-\frac{1}{4}\left(\Theta^{-1} \theta\right)\left(\Theta^{-1} \theta\right)^{T}}=\binom{\mu}{-\left(\Sigma+\mu \mu^{T}\right)} \\
\tilde{\Theta}=\nabla_{\tilde{H}} F^{*}(\tilde{H})=\binom{\nabla_{\tilde{H}} F^{*}(\eta)}{\nabla_{\tilde{H}} F^{*}(H)}=\binom{-\left(H+\eta \eta^{T}\right)^{-1} \eta}{-\frac{1}{2}\left(H+\eta \eta^{T}\right)^{-1}}=\binom{\Sigma^{-1} \mu}{\frac{1}{2} \Sigma^{-1}}
\end{gathered}
$$

## Example of entropic multivariate normal means

Simplify and extend (Myrvoll \& Soong'03)


Right in red Left in blue Symmetrized in green Geodesic half point $\lambda=\frac{1}{2}$ in purple

Left $D_{F}$ centroid is a right Kullback-Leibler centroid.
$\rightarrow$ Generalize the approach of (Davis \& Dhillon, NIPS'06).

## Clustering multivariate normal means

Center-based clustering (Banerjee et al., JMLR’05), (Teboulle'07)


Primal (natural ẽ)


Dual (expectation $\tilde{H})$

## COSH distance: Symmetrized Itakura-Saito

Burg entropy $F(x)=-\sum_{i=1}^{d} \log x^{(i)}$ yields Itakura-Saito div.

$$
\operatorname{IS}(p \| q)=\sum_{i=1}^{d}\left(\frac{p_{i}}{q_{i}}+\log \frac{p_{i}}{q_{i}}-1\right)=D_{F}(p \| q) .
$$

Symmetrized Itakura-Saito divergence is called COSH distance

$$
\operatorname{CosH}(p ; q)=\frac{\operatorname{IS}(p \| q)+\operatorname{IS}(q \| p)}{2}
$$

(Wei \& Gibson'00) COSH performs "best" for sound processing.

## Bregman-Csiszár centroids

(Csiszár Ann. Stat.'91) and (Lafferty COLT'99) used the following continuous family of generators:

$$
F_{\alpha}= \begin{cases}x-\log x-1 & \alpha=0 \\ \frac{1}{\alpha(1-\alpha)}\left(-x^{\alpha}+\alpha x-\alpha+1\right) & \alpha \in(0,1) \\ x \log x-x+1 & \alpha=1\end{cases}
$$

Continuum of Bregman divergences:

$$
\begin{aligned}
\mathrm{BC}_{0}(p \| q) & =\log \frac{q}{p}+\frac{p}{q}-1 \quad \text { Itakura-Saito } \\
\mathrm{BC}_{\alpha}(p \| q) & =\frac{1}{\alpha(1-\alpha)}\left(q^{\alpha}(x)-p^{\alpha}(x)+\alpha q(x)^{\alpha-1}(p(x)-q(x))\right)
\end{aligned}
$$

$\mathrm{BC}_{1}(p \| q)=p \log \frac{p}{q}-p+q$ Ext. Kullback-Leibler.
$\rightarrow$ Smooth spectrum of symmetric entropic centroids.

## Summary of presentation and contributions

Bregman centroids (Kullback-Leibler exp. fam. centroids)

- Left-sided and right-sided Bregman barycenters are generalized means,
- Symmetrized Bregman centroid $c^{F}$ is exactly geometrically characterized,
- Simple dichotomic geodesic walk to approximate $c^{F}$ (2-mean on sided centroids)
Generalize ad-hoc solutions with geometric interpretation:
- SKL centroid for non-parameter histograms (Veldhuis'02)
- SKL centroid for parametric multivariate normals (Myrvoll \& Soong'03)
- Right KL for multivariate normal (Davis \& Dhillon'06)


## Thank you!



References:
arXiv.org - cs - cs.CG (Computational Geometry) arXiv:0711.3242
http://www.sonycsl.co.jp/person/nielsen/BregmanCentroids/

## Acknowledgments:

Sony CSL, École Polytechnique (LIX), CEREGMIA, ANR GAIA.

