

The knee-jerk mapping

Peter G. Doyle Jim Reeds

Version dated 5 October 1998
GNU FDL*

Abstract

We claim to give the definitive theory of what we call the ‘knee-jerk mapping’, which is the basis for a class of optimization algorithms introduced by Baum, and promoted by Dempster, Laird, and Rubin under the name ‘EM algorithm’.

Introduction

We give the definitive theory of the knee-jerk mapping, to be defined below. This mapping has been investigated by many people, most notably Baum ([2], [3], [5], [4], [1]).

We begin with an example, taken from [6]. Suppose you want to locate the maximum of the function

$$Z(x, y) = x^{34}y^{38}(1 + 2x)^{125}$$

on the *1-simplex* (a fancy name for a line segment)

$$\Sigma = \{x, y > 0; x + y = 1\}.$$

*Copyright (C) 1990, 1998 Peter G. Doyle. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, as published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.

One way you can find it is by iterating the *knee-jerk mapping*

$$(x, y) \mapsto \frac{1}{xZ_x + yZ_y}(xZ_x, yZ_y) = \dots$$

This maps the simplex Σ to itself, and what is notable about the mapping is that it increases the value of the objective function Z .

The *one true explanation* of this ratcheting property of the knee-jerk map, the explanation that *lays bare* once and for all what is going on here, is as follows: Like any polynomial with only positive coefficients, the function Z is *log-log-convex*; that is, $\log Z$ is convex as a function of $(\log x, \log y)$; that is,

$$W(u, v) = \log Z(e^u, e^v)$$

is convex as a function of (u, v) . We're trying to find the maximum of W on the set

$$T = \{e^u + e^v = 1\}.$$

Since W is convex, if we fix a point (u, v) , the graph of W lies above its tangent plane at $(u, v, W(u, v))$:

$$W(\bar{u}, \bar{v}) \geq W_u(u, v)(\bar{u} - u) + W_v(u, v)(\bar{v} - v).$$

Now ideally we'd like to move from (u, v) directly to the point of T where $W(u, v)$ is greatest. What the knee-jerk mapping does is move instead to the point where the lower bound on the right hand side of the inequality above is maximized. This can't help increasing the objective function, right?

One remarkable fact should be pointed out, though it won't be gone into below: While the function Z is log-log-convex, it is nevertheless *log-concave*; that is, $\log Z$ is concave as a function of (x, y) . (This is true because Z is a product of homogeneous linear functions with positive coefficients.) Because Z is log-concave, it has a unique maximum on the simplex Σ . While all polynomials with positive coefficients are log-log-convex, only very special polynomials are simultaneously log-concave.

A class of log-concave examples fundamentally more exciting than products of linear functions can be obtained as follows: Take a connected graph G , think of its edges as variables, form for each spanning tree of G a monomial (of degree one smaller than the number of vertices of G), and form a

polynomial D_G —the *discriminant* of G —by adding up the monomials corresponding to all spanning trees of G . For example, if G is a triangle with edges x, y, z ,

$$D_G(x, y, z) = xy + xz + yz.$$

Discriminants of graphs are always log-concave. (If you know what a matroid is, let me add that the discriminant of a regular matroid is log-concave, but I don't know if the discriminant of a general matroid always is; my guess is that it isn't.)

Discriminants of graphs are particular cases of the *diagonal discriminants* of Bott and Duffin; these are always log-concave (because the determinant function is log-concave when restricted to the set of positive-definite matrices) as well as being log-log-convex (because they are polynomials with positive coefficients).

Knee-jerk functions

In real n -space, we will denote the positive orthant by Π and the closed standard simplex by Σ :

$$\Pi = \{x_1, \dots, x_n > 0\},$$

$$\Sigma = \{x_1, \dots, x_n > 0; x_1 + \dots + x_n = 1\}.$$

We denote their closures by $\bar{\Pi}$ (the non-negative orthant) and $\bar{\Sigma}$ (the closed standard simplex).

We say that a function $Z(x_1, \dots, x_n)$ from Π to the positive real numbers is *log-log-convex* if $\log Z$ is a convex function of $u_1 = \log x_1, \dots, u_n = \log x_n$. The name comes from the fact that in the case $n = 1$ a log-log-convex function is one whose graph appears convex when drawn on log-log graph paper. We say that Z is a *knee-jerk function* if Z is increasing (which we take to mean what some would call 'non-decreasing') and log-log-convex. For pedantry's sake we require in addition that Z be smooth, and extend continuously to $\bar{\Pi}$.

Properties and examples.

There are many characterizations of convex functions, but for our purposes the most important is that a function is convex if and only if its graph lies

above all of its tangent planes. Thus a smooth function Z is log-log-convex if and only if for any two points $\mathbf{x} = (x_1, \dots, x_n)$ and $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_n)$,

$$\begin{aligned} \log \bar{Z} - \log Z &\geq (\log Z)_{u_1}(\bar{u}_1 - u_1) + \dots + (\log Z)_{u_n}(\bar{u}_n - u_n) \\ &= \frac{x_1 Z_{x_1}}{Z} \log \frac{\bar{x}_1}{x_1} + \dots + \frac{x_n Z_{x_n}}{Z} \log \frac{\bar{x}_n}{x_n}, \end{aligned}$$

where $\bar{Z} = Z(\bar{\mathbf{x}})$ and $(\log Z)_{u_1}$ denotes the derivative of $\log Z$ with respect to u_1 , etc.

Using this characterization of log-log-convexity and Jensen's inequality—which states that for a concave function like \log the weighted average of the values is littler than the value of the weighted average—we get a proof that the function $Z(x_1, \dots, x_n) = x_1 + \dots + x_n$ is log-log-convex, and hence a knee-jerk function:

$$\begin{aligned} &\frac{x_1 Z_{x_1}}{Z} \log \frac{\bar{x}_1}{x_1} + \dots + \frac{x_n Z_{x_n}}{Z} \log \frac{\bar{x}_n}{x_n} \\ &= \frac{x_1}{x_1 + \dots + x_n} \log \frac{\bar{x}_1}{x_1} + \dots + \frac{x_n}{x_1 + \dots + x_n} \log \frac{\bar{x}_n}{x_n} \\ &\leq \log \left(\frac{x_1}{x_1 + \dots + x_n} \frac{\bar{x}_1}{x_1} + \dots + \frac{x_n}{x_1 + \dots + x_n} \frac{\bar{x}_n}{x_n} \right) \\ &= \log \frac{\bar{x}_1 + \dots + \bar{x}_n}{x_1 + \dots + x_n} \\ &= \log \frac{\bar{Z}}{Z}. \end{aligned}$$

Once we know that $x_1 + \dots + x_n$ is a knee-jerk function, we can easily produce a wealth of other examples by observing that the class of knee-jerk functions is closed under a variety of operations. The coordinate functions x_1, \dots, x_n are knee-jerk functions, as is any positive constant function. Products, positive scalar multiples, and positive (possibly fractional) powers of knee-jerk functions are knee-jerk functions. So is the composition $Z(Z_1, \dots, Z_k)$ of a knee-jerk function $Z(x_1, \dots, x_k)$ with knee-jerk functions $Z_1(x_1, \dots, x_n), \dots, Z_k(x_1, \dots, x_n)$, because the composition of increasing convex functions is increasing and convex. And since $x_1 + \dots + x_n$ is a knee-jerk function, it follows that sums of knee-jerk functions are knee-jerk functions. Thus any non-zero polynomial with non-negative coefficients is a knee-jerk function.

The knee-jerk mapping

If $Z(x_1, \dots, x_n)$ is a knee-jerk function, we define the *knee-jerk mapping*

$$\begin{aligned} T_Z(\mathbf{x}) &= \frac{1}{x_1 Z_{x_1} + \dots + x_n Z_{x_n}}(x_1 Z_{x_1}, \dots, x_n Z_{x_n}) \\ &= \frac{Z}{x_1 Z_{x_1} + \dots + x_n Z_{x_n}}(x_1 (\log Z)_{x_1}, \dots, x_n (\log Z)_{x_n}) \\ &= \frac{1}{(\log Z)_{x_1} + \dots + (\log Z)_{x_n}}((\log Z)_{x_1}, \dots, (\log Z)_{x_n}). \end{aligned}$$

(If $Z_{x_1} = \dots = Z_{x_n} = 0$, we define $T_Z(x_1, \dots, x_n) = \frac{1}{x_1 + \dots + x_n}(x_1, \dots, x_n)$ —or just pretend we didn't notice.) Note that when Z is homogeneous of (possibly fractional) degree d , Euler's identity

$$x_1 Z_{x_1} + \dots + x_n Z_{x_n} = dZ$$

implies that

$$T_Z(\mathbf{x}) = \frac{1}{dZ}(x_1 Z_{x_1}, \dots, x_n Z_{x_n}).$$

T_Z maps the positive orthant Π to the closed simplex $\bar{\Sigma}$, and thus restricts to a mapping of Σ to $\bar{\Sigma}$. It is easy to see that a point $\mathbf{x} \in \Sigma$ is fixed by T_Z if and only if it is a critical point of Z on Σ . The great thing about the knee-jerk mapping is that if \mathbf{x} is not a critical point of Z on Σ then $Z(T_Z(\mathbf{x})) > Z$; this will be proven in the next section. This makes the knee-jerk mapping a natural to iterate if you are interested in finding the maximum of Z on Σ . The name 'knee-jerk' is partly meant to suggest the automatic way in which the mapping increases the objective function Z .

The knee-jerk inequality

Write

$$\mathbf{x}' = T_Z(\mathbf{x})$$

and

$$Z' = Z(\mathbf{x}').$$

The knee-jerk inequality.

$$\log \frac{Z'}{Z} \geq \frac{x_1 Z_{x_1} + \dots + x_n Z_{x_n}}{Z} \left(x'_1 \log \frac{x'_1}{x_1} + \dots + x'_n \log \frac{x'_n}{x_n} \right).$$

Proof. From the characterization of log-log-convexity above, we have

$$\log \bar{Z} - \log Z \geq \frac{x_1 Z_{x_1}}{Z} \log \frac{\bar{x}_1}{x_1} + \dots + \frac{x_n Z_{x_n}}{Z} \log \frac{\bar{x}_n}{x_n}.$$

Substituting $\bar{\mathbf{x}} = \mathbf{x}'$ yields the knee-jerk inequality. ♠

Recall (if you don't already know) that for probability vectors $\mathbf{x} \in \Sigma$, $\mathbf{y} \in \bar{\Sigma}$ the *I-divergence* $I(\mathbf{y}; \mathbf{x})$ is defined to be

$$I(\mathbf{y}; \mathbf{x}) = y_1 \log \frac{y_1}{x_1} + \dots + y_n \log \frac{y_n}{x_n}.$$

This quantity is always ≥ 0 , with equality if and only if $\mathbf{x} = \mathbf{y}$. (This follows from an application of Jensen's inequality similar to that used above to show that $x_1 + \dots + x_n$ is a knee-jerk function.)

Corollary. *If $\mathbf{x} \in \Sigma$ then*

$$\log \frac{Z'}{Z} \geq \frac{x_1 Z_{x_1} + \dots + x_n Z_{x_n}}{Z} I(\mathbf{x}'; \mathbf{x}) \geq 0.$$

In particular, $Z' > Z$ unless the point \mathbf{x} is fixed by T_Z , which happens if and only if \mathbf{x} is a critical point of Z on Σ . ♠

What is going on here?

Say our goal is to maximize Z over Σ . We're sitting at some point \mathbf{x} , and we want to pick a new point $\bar{\mathbf{x}} \in \bar{\Sigma}$ so as to increase the objective function Z as much as possible. Since Z is log-log-convex we know that

$$\log \bar{Z} - \log Z \geq (\log Z)_{u_1} (\bar{u}_1 - u_1) + \dots + (\log Z)_{u_n} (\bar{u}_n - u_n).$$

The knee-jerk idea is to choose $\bar{\mathbf{x}} \in \bar{\Sigma}$ so as to make the lower bound on the right of this inequality as large as possible. That is, we want to do as well as

possible using only the value of Z and its derivatives at \mathbf{x} and the knowledge that Z is a knee-jerk function. So we want to choose $\bar{\mathbf{x}}$ so as to maximize

$$F(\bar{u}_1, \dots, \bar{u}_n) \equiv (\log Z)_{u_1}(\bar{u}_1 - u_1) + \dots + (\log Z)_{u_n}(\bar{u}_n - u_n)$$

subject to the constraint

$$G(\bar{u}_1, \dots, \bar{u}_n) \equiv e^{\bar{u}_1} + \dots + e^{\bar{u}_n} = 1.$$

The maximum occurs where

$$\nabla_{\bar{\mathbf{u}}} G = (\bar{x}_1, \dots, \bar{x}_n)$$

is proportional to

$$\nabla_{\bar{\mathbf{u}}} F = ((\log Z)_{u_1}, \dots, (\log Z)_{u_n}),$$

that is, where

$$\bar{\mathbf{x}} = T_Z(\mathbf{x}).$$

Ruminations. When $\mathbf{x} \in \Sigma$, the fact that $\mathbf{x}' = T_Z(\mathbf{x})$ maximizes the lower bound for \bar{Z} implies right away that $Z' \geq Z$, independently of the hocus-pocus with the I-divergence. Indeed, the positivity of the I-divergence can now be seen as a consequence of the fact that $x_1 + \dots + x_n$ is a knee-jerk function. This is not so surprising, perhaps, since both facts followed from very similar applications of Jensen's inequality. But now it appears that $x_1 + \dots + x_n$ is somehow the most important of all knee-jerk functions. And why should it be so distinguished? Because it crops up in the definition of the simplex Σ .

Generalizations

Given $\mathbf{a} = (a_1, \dots, a_n)$, $a_1, \dots, a_n > 0$, define

$$\Sigma_{\mathbf{a}} = \{x_1, \dots, x_n > 0; a_1 x_1 + \dots + a_n x_n = 1\}$$

and define

$$T_{Z, \mathbf{a}} : \Pi \rightarrow \bar{\Sigma}_{\mathbf{a}},$$

$$\mathbf{x}' = T_{Z,\mathbf{a}}(\mathbf{x}) = \frac{1}{x_1 Z_{x_1} + \dots + x_n Z_{x_n}} \left(\frac{x_1 Z_{x_1}}{a_1}, \dots, \frac{x_n Z_{x_n}}{a_n} \right).$$

Then the knee-jerk inequality becomes

$$\log \frac{Z'}{Z} \geq \frac{x_1 Z_{x_1} + \dots + x_n Z_{x_n}}{Z} \left(a_1 x'_1 \log \frac{x'_1}{x_1} + \dots + a_n x'_n \log \frac{x'_n}{x_n} \right).$$

When $\mathbf{x} \in \Sigma$ this becomes

$$\log \frac{Z'}{Z} \geq \frac{x_1 Z_{x_1} + \dots + x_n Z_{x_n}}{Z} I((a_1 x'_1, \dots, a_n x'_n); (a_1 x_1, \dots, a_n x_n)) \geq 0.$$

More interesting, we can replace the simplex Σ with a product of simplices: Let

$$Z = Z(x_{1,1}, \dots, x_{1,n_1}, \dots, x_{k,1}, \dots, x_{k,n_k}).$$

Let

$$T = \{x_{i,j} > 0; \sum_j x_{i,j} = 1\},$$

and define

$$T_Z : \Pi \mapsto \bar{T}$$

by

$$x'_{i,j} = \frac{x_{i,j} Z_{x_{i,j}}}{\sum_j x_{i,j} Z_{x_{i,j}}}.$$

Then

$$\log \frac{Z'}{Z} \geq \sum_i \frac{\sum_j x_{i,j} Z_{x_{i,j}}}{Z} \left(\sum_j x'_{i,j} \log \frac{x'_{i,j}}{x_{i,j}} \right),$$

and when $\mathbf{x} \in T$,

$$\log \frac{Z'}{Z} \geq \sum_i \frac{\sum_j x_{i,j} Z_{x_{i,j}}}{Z} I((x'_{i,1}, \dots, x'_{i,n_i}); (x_{i,1}, \dots, x_{i,n_i})) \geq 0.$$

References

- [1] L. E. Baum. An inequality and associated maximization technique in statistical estimation for probabilistic functions of Markov processes. In *Inequalities, Vol. 3*, pages 1–8. Academic Press, New York, 1972.

- [2] L. E. Baum and J. A. Eagon. An inequality with applications to statistical estimation for probabilistic functions of Markov processes and to a model for ecology. *Bull. Amer. Math. Soc.*, 73:360–363, 1967.
- [3] L. E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *Ann. Math. Stat.*, 37:1554–1563, 1966.
- [4] L. E. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Stat.*, 41:164–171, 1970.
- [5] L. E. Baum and G. R. Sell. Growth transformations for functions on manifolds. *Pacific J. Math.*, 27:211–227, 1968.
- [6] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Ann. Math. Stat.*, 41:164–171, 1970.