A Mathematical View of Interior-Point Methods for Convex Optimization

James Renegar[1]
School of Operations Research
and Industrial Engineering
Cornell University
renegar@orie.cornell.edu

July 28, 1998

[1]Supported by NSF Grant CCR-9403580

These lecture notes aim at developing a thorough understanding of the core theory for interior-point methods. The overall theory continues to grow at a rapid rate but the core ideas have remained largely unchanged for several years, since Nesterov and Nemirovskii[1] published their path-breaking, broadly-encompassing and enormously influential work. Since then, what has changed about the core ideas is our conception of them. Whereas [1] is notoriously difficult reading even for specialists, we now know how to motivate and present the general theory in such a way as to make it accessible for non-specialists and PhD students. Therein lies the justification for these lecture notes.

We develop the theory in $\Re^n$ although most of the theory can be developed in arbitrary real Hilbert spaces. The restriction to finite dimensions is primarily for accessibility.

The notes were developed largely in conjunction with a PhD-level course on interior-point methods at Cornell University.

Presently, the notes contain only two chapters, but those are sufficient to provide the reader with a solid introduction to the contemporary view of interior-point methods. A chapter on Duality Theory is nearing completion. A chapter on Complexity Theory is planned. If you are interested in receiving the chapters as they are completed, please send a brief message to renegar@orie.cornell.edu.

Chapter 1

Preliminaries

This chapter provides a review of material pertinent to continuous optimization theory quite generally, albeit phrased so as to be readily applicable in developing interior-point method (ipm) theory. The primary difference between our exposition and more customary approaches is that we do not rely on coordinate systems. For example, it is customary to define the gradient of a functional $f : \Re^n \to \Re$ as the vector-valued function $g : \Re^n \to \Re^n$ whose j^{th} coordinate is $\partial f / \partial x_j$. Instead, we consider the gradient as determined by an underlying inner product $\langle \, , \, \rangle$. For us, the gradient is the function g satisfying

$$\lim_{\|\Delta x\| \to 0} \frac{f(x + \Delta x) - f(x) - \langle g(x), \Delta x \rangle}{\|\Delta x\|} = 0,$$

where $\|\Delta x\| := \langle \Delta x, \Delta x \rangle^{1/2}$. In general, the function whose j^{th} coordinate is $\partial f / \partial x_j$ is the gradient only if $\langle \, , \, \rangle$ is the Euclidean inner product.

The natural geometry varies from point to point in the domains of optimization problems that can be solved by ipm's. As the algorithms progress from one point to the next, one changes the inner product – and hence the geometry – to visualize the headway achieved by the algorithms. The relevant inner products may bear no relation to an initially-imposed coordinate-system. Consequently, in aiming for the most transparent and least cumbersome proofs, one should dispense with coordinate systems.

We begin with a review of linear algebra by recalling, for example, the notion of a self-adjoint linear operator. We then define gradients and Hessians, emphasizing how they change when the underlying inner product is changed. Next is a brief review of basic results for convex functionals, followed by results akin to the Fundamental Theorem of Calculus. Although these "Calculus results" are elementary and rather dry, they are essential in achieving lift-off for the ipm theory. Finally, we recall Newton's method for continuous optimization, proving a standard theorem which later plays a central motivational role.

1.1 Linear Algebra

We let $\langle\ ,\ \rangle$ denote an arbitrary inner product on $\Re^n$. In later sections, $\langle\ ,\ \rangle$ will act as a reference inner product, an inner product from which other inner products are constructed. For the ipm theory, it happens that the reference inner product $\langle\ ,\ \rangle$ is irrelevant; the inner products essential to the theory are independent of the reference inner product. To large extent, the reference inner product will serve only to fix notation.

Although the particular reference inner product will prove to be irrelevant for ipm theory, for optimization problems to be solved by ipm's there typically are associated natural reference inner products. For example, in linear programming (LP) where vectors x are expressed coordinate-wise and "$x \geq 0$" means each coordinate is non-negative, the natural inner product is the Euclidean inner product, which we refer to as the "dot product," writing $x_1 \cdot x_2$. Similarly, in semi-definite programming (SDP) where the relevant vector space is $\mathcal{S}^{n \times n}$ – the space of symmetric $n \times n$ real matrices X – and "$X \succeq 0$" means X is positive semi-definite (i.e., has no negative eigenvalues), the natural inner product is the "trace product,"

$$X \circ S := \mathrm{trace}(XS).$$

(Thus, $X \circ S$ equals the sum of the eigenvalues of the matrix XS.)

Throughout the general development, we use $\Re^n$ to denote an arbitrary finite-dimensional real vector space, be it $\mathcal{S}^{n \times n}$ or whatever.

The inner product $\langle\ ,\ \rangle$ induces a norm on $\Re^n$,

$$\|x\| := \langle x, x\rangle^{1/2}.$$

Perhaps the most useful relation between the inner product and the norm is the *Cauchy-Schwarz inequality,*

$$|\langle x_1, x_2\rangle| \leq \|x_1\|\,\|x_2\|$$

with equality iff x_1 and x_2 are co-linear. If neither x_1 nor x_2 is the zero vector, Cauchy-Schwarz implies the existence of Θ satisfying

$$\cos\Theta = \langle x_1, x_2\rangle / \|x_1\|\,\|x_2\|.$$

The value Θ is referred to as the angle between x_1 and x_2.

Whereas the dot product gives rise to the Euclidean norm, the norm arising from the trace product is known as the "Frobenius norm." The Frobenius norm can be extended to the vector space of all real $n \times n$ matrices by defining $\|M\| := (\sum m_{ij}^2)^{1/2}$ where m_{ij} are the coefficients of M.

We remark that the Frobenius norm is "submultiplicative," meaning $\|XS\| \leq \|X\| \, \|S\|$.

Recall that vectors $x_1, x_2 \in \Re^n$ are said to be *orthogonal* if $\langle x_1, x_2 \rangle = 0$. Recall that a basis $v_1, \ldots, v_n$ for $\Re^n$ is said to be an *orthonormal* if

$$\langle v_i, v_j \rangle = \delta_{ij} \text{ for all } i, j$$

where δ_{ij} is the Kronecker delta. A linear operator (i.e., a linear transformation) $Q : \Re^n \to \Re^n$ is said to be *orthogonal* if

$$\langle Qx_1, Qx_2 \rangle = \langle x_1, x_2 \rangle \text{ for all } x_1, x_2 \in \Re^n.$$

If given an inner product $\langle \; , \; \rangle$ on $\Re^n$, one adopts a coordinate system obtained by expressing vectors as linear combinations of an orthonormal basis, the inner product $\langle \; , \; \rangle$ is the dot product for that coordinate system. Consequently, one can consider the results we review below as following from the special case of the dot product, but one should keep in mind that thinking in terms of coordinates is best avoided for understanding the ipm theory.

If both $\Re^n$ and $\Re^m$ are endowed with inner products and $A : \Re^n \to \Re^m$ is a linear operator, there exists a unique linear operator $A^* : \Re^m \to \Re^n$ satisfying

$$\langle Ax, y \rangle = \langle x, A^* y \rangle \text{ for all } x \in \Re^n, \, y \in \Re^m.$$

The operator A^* is the *adjoint* of A. The range of A^* is orthogonal to the nullspace of A.

Assuming A is surjective, the linear operator $A^*(AA^*)^{-1}A$ projects $\Re^n$ orthogonally onto the range space of A^*, that is, the image of x is the point in the range space closest to x. Likewise, $I - A^*(AA^*)^{-1}A$ projects $\Re^n$ orthogonally onto the nullspace of A.

If both $\Re^n$ and $\Re^m$ are endowed with the dot product and if A and A^* are written as matrices then A^* is the transpose of A. Thus it is natural in this setting to write A^T rather than A^*.

It is a simple but important exercise for SDP to show that if $S_1, \ldots, S_m \in \mathcal{S}^{n \times n}$ and $A : \mathcal{S}^{n \times n} \to \Re^m$ is the linear operator defined by

$$X \mapsto (X \circ S_1, \ldots, X \circ S_m)$$

then

$$A^*y = \sum_i y_i S_i,$$

assuming $\mathcal{S}^{n\times n}$ is endowed with the trace product and $\Re^m$ is endowed with the dot product.

Continuing to assume $\Re^n$ and $\Re^m$ are endowed with inner products, and hence norms, one obtains an induced *operator norm* on the vector space consisting of linear operators $A : \Re^n \to \Re^m$:

$$\|A\| := \max\{\|Ax\| : \|x\| \leq 1\}.$$

Each linear operator $A : \Re^n \to \Re^m$ has a *singular-value decomposition.* Precisely, there exist orthomormal bases $u_1, \ldots, u_n$ and $w_1, \ldots, w_m$, as well as real numbers $0 < \gamma_1 \leq \ldots \leq \gamma_r$ where r is the rank of A, such that for all x,

$$Ax = \sum_{i=1}^{r} \gamma_i \langle u_i, x\rangle w_i.$$

The numbers γ_i are the *singular values* of A; if $r < n$ then the number 0 is also considered to be a singular value of A. It is easily seen that $\|A\| = \gamma_r$. Moreover,

$$A^*y = \sum_{i=1}^{r} \gamma_i \langle w_i, y\rangle u_i,$$

so that the values γ_i (and possibly 0) are also the singular values of A^*. It immediately follows that $\|A^*\| = \|A\|$.

If $\Re^n$ and $\Re^m$ are endowed with the dot product, the singular-value decomposition corresponds to the fact that if A is an $m \times n$ matrix, there exist orthogonal matrices Q_m and Q_n such that $Q_m A Q_n = \Gamma$ where Γ is an $m \times n$ matrix with zeros everywhere except possibly for positive numbers on its main diagonal.

It is not difficult to prove that a linear operator $Q : \Re^n \to \Re^n$ is orthogonal iff $Q^* = Q^{-1}$. For orthogonal operators, $\|Q\| = 1$.

A linear operator $S : \Re^n \to \Re^n$ is said to be *self-adjoint* if $S = S^*$.

If $\langle \ , \ \rangle$ is the dot product and S is written as a matrix then S being self-adjoint is equivalent to S being symmetric.

It is instructive to show that for $S \in \mathcal{S}^{n\times n}$, the linear operator $A : \mathcal{S}^{n\times n} \to \mathcal{S}^{n\times n}$ defined by

$$X \mapsto SXS$$

is self-adjoint. Such operators are important in the ipm theory for SDP.

A linear operator $S : \Re^n \to \Re^n$ is said to be *positive semi-definite* (psd) if S is self-adjoint and

$$\langle x, Sx \rangle \geq 0 \text{ for all } x \in \Re^n.$$

If, further, S satisfies

$$\langle x, Sx \rangle > 0 \text{ for all } x \neq 0$$

then S is said to be *positive definite* (pd).

Each self-adjoint linear operator S has a *spectral decomposition.* Precisely, for each self-adjoint linear operator S there exists an orthonormal basis $v_1, \ldots, v_n$ and real numbers $\lambda_1 \leq \ldots \leq \lambda_n$ such that for all x,

$$Sx = \sum \lambda_i \langle v_i, x \rangle v_i.$$

It is easily seen that v_i is an eigenvector for S with eigenvalue λ_i.

If $\langle \ , \ \rangle$ is the dot product and S is a symmetric matrix then the spectral decomposition corresponds to the fact that S can be diagonalized using an orthogonal matrix, i.e., $Q^T S Q = \Lambda$.

The following relations are easily established:

- $\|S\| = \max_i |\lambda_i| = \max\{|\langle x, Sx \rangle| : \|x\| = 1\}$;
- S is psd iff $\lambda_i \geq 0$ for all i;
- S is pd iff $\lambda_i > 0$ for all i;
- If S^{-1} exists then it, too, is self-adjoint, and has eigenvalues $1/\lambda_i$. (In particular, $\|S^{-1}\| = 1/\min_i |\lambda_i|$.)

The spectral decomposition for a psd operator S allows one to easily prove the existence of a psd operator $S^{1/2}$ satisfying $S = (S^{1/2})^2$; simply replace λ_i by $\sqrt{\lambda_i}$ in the decomposition. In turn, the uniqueness of $S^{1/2}$ can readily be proven by relying on the fact that if T is a psd operator satisfying $T^2 = S$ then the eigenvectors for T are eigenvectors for S. The operator $S^{1/2}$ is the "square root" of S.

Here is a crucial observation: If S is pd then S defines a new inner product, namely,

$$\langle x_1, x_2 \rangle_S := \langle x_1, S x_2 \rangle.$$

Every inner product on $\Re^n$ arises in this way; that is, regardless of the initial inner product $\langle \ , \ \rangle$, for every other inner product there exists S

which is pd w.r.t. $\langle\ ,\ \rangle$ and for which $\langle\ ,\ \rangle_S$ is precisely the other inner product.

Let $\|\ \|_S$ denote the norm induced by $\langle\ ,\ \rangle_S$.

Assume A^* is the adjoint of $A : \Re^n \to \Re^m$. Assuming S and T are pd w.r.t. the respective inner products, if the inner product on $\Re^n$ is replaced by $\langle\ ,\ \rangle_S$ and that on $\Re^m$ is replaced by $\langle\ ,\ \rangle_T$ then the adjoint of A becomes $S^{-1}A^*T$, as is easily shown. In particular, if $m = n$ and $S = T$ then the adjoint of A becomes $S^{-1}A^*S$. Moreover, letting $\|A\|_{S,T}$ denote the resulting operator norm, it is easily proven that

$$\|A\|_{S,T} = \|T^{1/2}AS^{-1/2}\| \quad \text{and} \quad \|A\| = \|T^{-1/2}AS^{1/2}\|_{S,T}.$$

The notation used in stating these facts illustrates our earlier assertion that the reference inner product will be useful in fixing notation (as the inner products change).

It is instructive to consider the shape of the unit ball w.r.t. $\|\ \|_S$ viewed in terms of the geometry of the reference inner product. The spectral decomposition of S easily implies the unit ball to be an ellipsoid with axes in the directions of the orthonormal basis vectors $v_1, \ldots, v_n$, the length of the axis in the direction of v_i being $2\sqrt{1/\lambda_i}$.

1.2 Gradients

Recall that a *functional* is a function whose range lies in $\Re$. We use D_f to denote the domain of a functional f. It will always be assumed that D_f is an open subset of $\Re^n$ in the norm topology (recalling that all norms on $\Re^n$ induce the same topology). Let $\langle\ ,\ \rangle$ denote an arbitrary inner product on $\Re^n$ and let $\|\ \|$ denote the norm induced by $\langle\ ,\ \rangle$.

The functional f is said to be (Frechet) *differentiable* at $x \in D_f$ if there exists a vector $g(x)$ satisfying

$$\lim_{\|\Delta x\| \to 0} \frac{f(x + \Delta x) - f(x) - \langle g(x), \Delta x \rangle}{\|\Delta x\|} = 0.$$

The vector $g(x)$ is the *gradient* of f at x w.r.t. $\langle\ ,\ \rangle$.

Of course if one chooses the inner product $\langle\ ,\ \rangle$ on $\Re^n$ to be the dot product and expresses $g(x)$ coordinate-wise, the j^{th} coordinate is $\partial f / \partial x_j$.

For an arbitrary inner product $\langle\ ,\ \rangle$, the gradient has the same geometrical interpretation that is taught in Calculus for the dot product gradient. Roughly speaking, the gradient $g(x)$ points in the direction for which the functional output increases the fastest per unit distance travelled, and the magnitude $\|g(x)\|$ equals the amount the functional will change per unit distance travelled in that direction. We give rigour to this geometrical interpretation in §1.5.

The *first-order approximation of f at x* is the linear functional

$$y \mapsto f(x) + \langle g(x), y - x \rangle.$$

If f is differentiable at each $x \in D_f$ then f is said to be *differentiable.* Henceforth, assume f is differentiable.

If the function $x \mapsto g(x)$ is continuous at each $x \in D_f$ then f is said to be *continuously differentiable.* One then writes $f \in C^1$.

To illustrate the definition of gradient, consider the functional

$$f(X) := -\ln\det(X)$$

with domain $\mathcal{S}^{n\times n}_{++}$, the set of all pd matrices in $\mathcal{S}^{n\times n}$. (This functional plays an especially important role in SDP.) We claim that w.r.t. the trace product,

$$g(X) = -X^{-1}.$$

For let $\Delta X \in \mathcal{S}^{n\times n}$ and denote the eigenvalues of $X^{-1/2}(\Delta X)X^{-1/2}$ by $\gamma_1, \ldots, \gamma_n$. Since the trace of a matrix depends only on the eigenvalues – hence, $X^{-1} \circ \Delta X = \operatorname{trace}(X^{-1/2}(\Delta X)X^{-1/2})$ – we have

$$\begin{aligned} & \frac{f(X + \Delta X) - f(X) - \langle -X^{-1}, \Delta X\rangle}{\|\Delta X\|} \\ = \; & \frac{|-\ln\det(X + \Delta X) + \ln\det(X) + X^{-1} \circ \Delta X|}{(\Delta X \circ \Delta X)^{1/2}} \\ = \; & \frac{|-\ln\det(I + X^{-1/2}(\Delta X)X^{-1/2}) + \operatorname{trace}(X^{-1/2}(\Delta X)X^{-1/2})|}{\operatorname{trace}(X^2)^{1/2}} \\ = \; & \frac{|\sum_i \gamma_i - \ln(1+\gamma_i)|}{\operatorname{trace}((\Delta X)^2)^{1/2}}. \end{aligned}$$

Letting $\lambda_i(X)$ and $\lambda_i(\Delta X)$ denote the eigenvalues of X and ΔX, it is easily proven that

$$\text{(1.2.1)} \quad \operatorname{trace}((\Delta X)^2)^{1/2} \geq \max_i |\lambda_i(\Delta X)| \geq (\min_i \lambda_i(X))(\max_i |\gamma_i|)$$

and hence

$$\begin{aligned} & \limsup_{\|\Delta X\| \to 0} \frac{f(X + \Delta X) - f(X) - \langle -X^{-1}, \Delta X\rangle}{\|\Delta X\|} \\ \leq \; & \frac{1}{\min_i \lambda_i(X)} \limsup_{\|\Delta X\| \to 0} \frac{|\sum_i \gamma_i - \ln(1+\gamma_i)|}{\max_i |\gamma_i|}. \end{aligned}$$

Since (1.2.1) implies $\gamma_i \to 0$ when $\|\Delta X\| \to 0$, it is now straightforward to conclude that the value of the limit supremum is 0. Thus, $g(X) = -X^{-1}$.

Our definition of what it means for a functional f to be differentiable depends on the inner product $\langle \, , \, \rangle$. However, relying on the equivalence of all norm topologies on $\Re^n$, it is readily proven that the property of being differentiable – and being continuously differentiable – is independent of the inner product. The gradient depends on the inner product but differentiability does not.

The following proposition shows how the gradient changes as the inner product changes.

Proposition 1.2.2 *If S is pd and f is differentiable at x then the gradient of f at x w.r.t. $\langle \, , \, \rangle_S$ is $S^{-1}g(x)$.*

Proof. Letting λ_1 denote the least eigenvalue of S, the proof relies on the fact that

$$\sqrt{\lambda_1}\|x\| \leq \|x\|_S \text{ for all } x,$$

as follows easily from the spectral decomposition of S.

To prove that $S^{-1}g(x)$ is the gradient of f at x w.r.t. $\langle \, , \, \rangle_S$, we wish to show

$$\limsup_{\|\Delta x\|_S \to 0} \frac{|f(x+\Delta x) - f(x) - \langle S^{-1}g(x), \Delta x\rangle_S|}{\|\Delta x\|_S} = 0.$$

However, noting that for all v,

$$\langle S^{-1}g(x), v\rangle_S = \langle S^{-1}g(x), Sv\rangle = \langle SS^{-1}g(x), v\rangle = \langle g(x), v\rangle$$

we have

$$\begin{aligned}
&\limsup_{\|\Delta x\|_S \to 0} \frac{|f(x+\Delta x) - f(x) - \langle S^{-1}g(x), \Delta x\rangle_S|}{\|\Delta x\|_S} \\
&= \limsup_{\|\Delta x\|_S \to 0} \frac{|f(x+\Delta x) - f(x) - \langle g(x), \Delta x\rangle|}{\|\Delta x\|_S} \\
&\leq \frac{1}{\sqrt{\lambda_1}} \limsup_{\|\Delta x\|_S \to 0} \frac{|f(x+\Delta x) - f(x) - \langle g(x), \Delta x\rangle|}{\|\Delta x\|} \\
&= \frac{1}{\sqrt{\lambda_1}} \lim_{\|\Delta x\| \to 0} \frac{f(x+\Delta x) - f(x) - \langle g(x), \Delta x\rangle}{\|\Delta x\|} = 0,
\end{aligned}$$

the next-to-last equality due to $\|\Delta x\| \to 0$ if $\|\Delta x\|_S \to 0$ (because $\|\Delta x\| \leq \|\Delta x\|_S/\sqrt{\lambda_1}$). □

Theorem 1.2.2 has the unsurprising consequence that the first-order approximation of f at x is independent of the inner product:

$$f(x) + \langle g(x), y - x\rangle = f(x) + \langle S^{-1}g(x), y - x\rangle_S.$$

Finally, we make an observation that will be important for applying ipm theory to optimization problems having linear equations among the constraints. Assume L is a subspace of $\Re^n$. Restricting $\langle\ ,\ \rangle$ to L makes L into an inner product space. Thus, if f is a functional for which $D_f \cap L \neq \emptyset$, one can speak of the gradient of $f|_L$, the functional obtained by restricting f to L. Let $g|_L$ denote the gradient, a function from L to L. It is not difficult to prove $g|_L = P_L g$ where P_L is the operator projecting $\Re^n$ orthogonally onto L. Summarizing, if $x \in L$ then the gradient $g|_L(x)$ of $f|_L$ at x is the vector $P_L g(x)$.

1.3 Hessians

The functional f is said to be *twice differentiable* at $x \in D_f$ if $f \in C^1$ and there exists a linear operator $H(x) : \Re^n \to \Re^n$ satisfying

$$\lim_{\|\Delta x\| \to 0} \frac{\|g(x + \Delta x) - g(x) - H(x)\Delta x\|}{\|\Delta x\|} = 0.$$

If it exists, $H(x)$ is said to be the *Hessian* of f at x w.r.t. $\langle\ ,\ \rangle$.

If $\langle\ ,\ \rangle$ is the dot product and $H(x)$ is written as a matrix, the (i, j) entry of $H(x)$ is

$$\frac{\partial^2 f}{\partial x_i \partial x_j}.$$

The *second-order approximation of f at x* is the linear functional

$$y \mapsto f(x) + \langle g(x), y - x\rangle + \tfrac{1}{2}\langle (y - x), H(x)(y - x)\rangle.$$

If f is twice differentiable at each $x \in D_f$ then f is said to be *twice differentiable.* Henceforth, assume f is twice differentiable.

If the function $x \mapsto H(x)$ is continuous at x (w.r.t. the operator-norm topology, or equivalently, any norm topology on the vector space of linear operators from $\Re^n$ to $\Re^n$), then $H(x)$ is self-adjoint. If the function $x \mapsto H(x)$ is continuous at each $x \in D_f$ then f is said to be *twice continuously differentiable.* One then writes $f \in C^2$.

The assumption of twice continuous differentiability, as opposed to mere twice differentiability, is often made in optimization primarily to ensure self-adjointness of the Hessian.

If the inner product is the dot product and Hessian is expressed as a matrix, self-adjointness is equivalent to the matrix being symmetric, that is,

$$\frac{\partial^2 f}{\partial x_i \partial x_j} = \frac{\partial^2 f}{\partial x_j \partial x_i}.$$

(If the Hessian matrix does not vary continuously in x, the order in which the partials are taken can matter, resulting in a non-symmetric matrix.)

To illustrate the definition of the Hessian, we again consider the functional

$$f(X) := -\ln\det(X)$$

with domain $\mathcal{S}^{n\times n}_{++}$, the set of all pd matrices in $\mathcal{S}^{n\times n}$. We saw that $g(X) = -X^{-1}$. We claim that $H(X)$ is the linear operator given by

$$\Delta X \mapsto X^{-1}(\Delta X)X^{-1}.$$

This can be proven by relying on the fact that if $\|\Delta X\|$ is sufficiently small then

$$(X + \Delta X)^{-1} = X^{-1}\sum_{k=0}^{\infty}[-(\Delta X)X^{-1}]^k,$$

and hence

$$g(X + \Delta X) - g(X) - H(X)\Delta X = -X^{-1}\sum_{k=2}^{\infty}[-(\Delta X)X^{-1}]^k.$$

For then, from the submultiplicativity of the Frobenius norm,

$$\begin{aligned} & \limsup_{\|\Delta X\|\to 0} \frac{|g(X+\Delta X) - g(X) - H(X)\Delta X|}{\|\Delta X\|} \\ \leq \; & \limsup_{\|\Delta X\|\to 0} \frac{\|\Delta X\|^2 \|X^{-1}\|^3 \sum_{k=0}^{\infty}(\|(\Delta X)\|\,\|X^{-1}\|)^k}{\|\Delta X\|} \\ = \; & 0. \end{aligned}$$

The property of being twice continuously differentiable does not depend on the inner product whereas the Hessian most certainly does depend on the inner product, as is made explicit in the following proposition. The proof of the following proposition is similar to the proof of Proposition 1.2.2 and hence is left to the reader.

Proposition 1.3.1 *If S is pd and f is twice differentiable at x then the Hessian of f at x w.r.t. $\langle\ ,\ \rangle_S$ is $S^{-1}H(x)$.*

Propositions 1.2.2 and 1.3.1 have the unsurprising consequence that the second-order approximation of f at x is independent of the inner product:

$$\begin{aligned} &f(x) + \langle g(x), y-x\rangle + \tfrac{1}{2}\langle (y-x), H(x)(y-x)\rangle \\ &= f(x) + \langle S^{-1}g(x), y-x\rangle_S + \tfrac{1}{2}\langle (y-x), S^{-1}H(x)(y-x)\rangle_S. \end{aligned}$$

Finally, we make an observation regarding Hessians and subspaces. It is straightforward to prove that if L is a subspace of $\Re^n$ and f satisfies $D_f \cap L \neq \emptyset$ then the Hessian of $f|_L$ at $x \in L$ – an operator from L to L – is given by $H|_L(x) = P_L H(x)$. That is, when one applies $H|_L(x)$ to a vector $v \in L$, one obtains the vector $P_L H(x)v$.

1.4 Convexity

Recall that a set $S \subseteq \Re^n$ is said to be *convex* if whenever $x, y \in S$ and $0 \leq t \leq 1$ we have $x + t(y-x) \in S$.

Recall that a functional f is said to be *convex* if D_f is convex and if whenever $x, y \in D_f$ and $0 \leq t \leq 1$, we have

$$f(x + t(y-x)) \leq f(x) + t(f(y) - f(x)).$$

If the inequality is strict whenever $0 < t < 1$ and $x \neq y$, then f is said to be *strictly convex.*

The minimizers of a convex functional form a convex set. A strictly convex functional has at most one minimizer.

Henceforth, we assume $f \in C^2$ and we assume D_f is an open, convex set.

If f is a univariate functional, we know from Calculus that f is convex iff $f''(x) \geq 0$ for all $x \in D_f$. Similarly, if $f''(x) > 0$ for all $x \in D_f$ then f is strictly convex. The following standard theorem generalizes these facts.

Proposition 1.4.1 *The functional f is convex iff $H(x)$ is psd for all $x \in D_f$. If $H(x)$ is pd for all $x \in D_f$ then f is strictly convex.*

The following elementary proposition, which is relied on in the proof of Proposition 1.4.1, is fundamental throughout these lecture notes. It does not assume convexity of f.

Proposition 1.4.2 *Assume $x, y \in D_f$ and define a univariate functional $\phi : [0,1] \to \Re$ by*

$$\phi(t) := f(x + t(y - x)).$$

Then

$$\phi'(t) = \langle g(x + t(y - x)), y - x \rangle$$

and

$$\phi''(t) = \langle y - x, H(x + t(y - x))(y - x) \rangle.$$

Proof. Fix t and let $u = x + t(y - x)$. We wish to prove

$$\phi'(t) = \langle g(u), y - x \rangle \text{ and } \phi''(t) = \langle y - x, H(u)(y - x) \rangle.$$

To prove $\phi'(t) = \langle g(u), y - x \rangle$ it suffices to show

$$\limsup_{s \to 0} \left| \frac{\phi(t + s) - \phi(t) - s\langle g(u), y - x \rangle}{s} \right| = 0.$$

However, noting $\phi(t) = f(u)$ and $\phi(t + s) = f(u + s(y - x))$, we have

$$\limsup_{s \to 0} \left| \frac{\phi(t + s) - \phi(t) - s\langle g(u), y - x \rangle}{s} \right|$$

$$= \limsup_{s \to 0} \left| \frac{f(u + s(y - x)) - f(u) - \langle g(u), s(y - x) \rangle}{s} \right|$$

$$= \|y - x\| \limsup_{\|s(y-x)\| \to 0} \frac{|f(u + s(y - x)) - f(u) - \langle g(u), s(y - x) \rangle|}{\|s(y - x)\|} = 0,$$

the final equality by definition of $g(u)$.

Similarly, to prove $\phi''(t) = \langle y - x, H(u)(y - x) \rangle$, it suffices to show

$$\limsup_{s \to 0} \left| \frac{\phi'(t + s) - \phi'(t) - s\langle y - x, H(u)(y - x) \rangle}{s} \right| = 0.$$

However, since we now know

$$\phi'(t) = \langle g(u), y - x \rangle \text{ and } \phi'(t + s) = \langle g(u + s(y - x)), y - x \rangle$$

we have

$$\limsup_{s \to 0} \left| \frac{\phi'(t + s) - \phi'(t) - s\langle y - x, H(u)(y - x) \rangle}{s} \right|$$

$$= \limsup_{s\to 0} \left| \frac{\langle g(u+s(y-x)) - g(u) - H(u)s(y-x), y-x\rangle}{s} \right|$$

$$\leq \limsup_{s\to 0} \frac{\|g(u+s(y-x)) - g(u) - H(u)s(y-x)\| \, \|y-x\|}{|s|}$$

$$= \|y-x\|^2 \limsup_{\|s(y-x)\|\to 0} \frac{\|g(u+s(y-x)) - g(u) - H(u)s(y-x)\|}{\|s(y-x)\|} = 0$$

where the inequality is by Cauchy-Schwarz and the final equality is by definition of $H(u)$. □

Proof of Proposition 1.4.1: We first show that if the Hessian is psd everywhere on D_f then f is convex. So assume the Hessian is psd everywhere on D_f.

Assume x and y are arbitrary points in D_f. We wish to show that if t satisfies $0 \leq t \leq 1$ then

$$f(x+t(y-x)) \leq f(x) + t(f(y) - f(x)). \tag{1.4.3}$$

Consider the univariate functional ϕ defined by

$$\phi(t) := f(x+t(y-x)).$$

Observe (1.4.3) is equivalent to

$$\phi(t) \leq \phi(0) + t(\phi(1) - \phi(0)),$$

an inquality that is certainly valid if ϕ is convex on the interval $[0,1]$. Hence to prove (1.4.3) it suffices to prove $\phi''(t) \geq 0$ for all $0 \leq t \leq 1$. However, Proposition 1.4.2 implies

$$\phi''(t) = \langle y-x, H(x+t(y-x))(y-x)\rangle \geq 0,$$

the inequality because $H(x+t(y-x))$ is psd.

The proof that f is strictly convex if the Hessian is everywhere pd on D_f is similar and hence is left to the reader.

To conclude the proof, it suffices to show that if $H(x)$ is not psd for some x then f is not convex. If $H(x)$ is not psd then $H(x)$ has an eigenvector v with negative eigenvalue λ. To show f is not convex, it suffices to show the functional $\phi(t) := f(x+tv)$ is not convex. To show ϕ is not convex, it suffices to show $\phi''(0) < 0$. This is straightforward, again relying on Proposition 1.4.2. □

Earlier in the notes it was asserted that, roughly speaking, the gradient $g(x)$ points in the direction for which the functional output increases the fastest per unit distance travelled, and the magnitude $\|g(x)\|$ equals the amount the functional will change per unit distance travelled in that direction. Proposition 1.4.2 provides the means to make this rigorous: Choose an arbitrary direction v of unit length. The initial rate of change in the output of f as one moves from x to $x+v$ in unit time is given by $\phi'(0)$ where

$$\phi(t) := f(x+tv).$$

Note Proposition 1.4.2 implies

$$\phi'(0) = \langle g(x), v\rangle \tag{1.4.4}$$

and hence, by Cauchy-Schwarz and $\|v\| = 1$, we have

$$-\|g(x)\| \leq \phi'(0) \leq \|g(x)\|.$$

So the initial rate of change cannot exceed $\|g(x)\|$ in magnitude, regardless of which direction v of unit length is chosen. However, assuming $g(x) \neq 0$, if one chooses the direction

$$v = \frac{1}{\|g(x)\|}g(x),$$

(1.4.4) implies $\phi'(0) = \|g(x)\|$.

We mention that a point z minimizes a convex functional f iff $g(z) = 0$. (For the "if," assume $g(z) = 0$. For $y \in D_f$, consider the univariate functional $\phi(t) := f(z + t(y-z))$. By Proposition 1.4.2, $\phi'(0) = 0$. Since ϕ is convex, we know from univariate Calculus that 0 minimizes ϕ. In particular, $\phi(0) \leq \phi(1)$, that is, $f(z) \leq f(y)$. For the "only if," assume $g(z) \neq 0$ and consider the discussion of the preceding paragraph.)

As with the two preceding sections, we close this one with a discussion of subspaces.

If L is a subspace of $\Re^n$ and $D_f \cap L \neq \emptyset$, we know the gradient of $f|_L$ to be $P_L g$. Thus, $z \in L$ solves the constrained optimization problem

$$\begin{array}{ll} \min & f(x) \\ \text{s.t.} & x \in L \end{array}$$

iff $P_L g(z) = 0$, that is, iff $g(z)$ is orthogonal to L. In particular, if L is the nullspace of a linear operator A then z solves the optimization problem iff $g(z) = A^* y$ for some y. Likewise when L is replaced by a translate of L; that is, when L is replaced by an affine space $v + L$ for some vector v. We record this in the following proposition.

Proposition 1.4.5 *If f is convex and A is a linear operator then $z \in D_f$ solves the linearly-constrained optimization problem*

$$\begin{array}{ll} \min & f(x) \\ s.t. & Ax = b \end{array}$$

iff $Az = b$ and $g(z) = A^*y$ for some y.

1.5 Fundamental Theorems of Calculus

We continue to assume $f \in C^2$ and D_f is an open, convex set.

The following theorem generalizes the Fundamental Theorem of Calculus.

Theorem 1.5.1 *If $x, y \in D_f$ then*

$$f(y) - f(x) = \int_0^1 \langle g(x + t(y-x)), y - x\rangle \, dt.$$

Proof. Consider the univariate functional $\phi(t) := f(x + t(y - x))$. The fundamental theorem of Calculus asserts

$$\phi(1) - \phi(0) = \int_0^1 \phi'(t)\, dt.$$

Since $\phi(1) = f(y)$, $\phi(0) = f(x)$ and, by Proposition 1.4.2,

$$\phi'(t) = \langle g(x + t(y - x)), y - x\rangle,$$

the proof is complete. □

In a similar vein, we have the following proposition.

Proposition 1.5.2 *If $x, y \in D_f$ then*

$$f(y) = f(x) + \langle g(x), y - x\rangle + \int_0^1 \langle g(x + t(y - x)) - g(x), y - x\rangle \, dt \tag{1.5.3}$$

and

$$f(y) = f(x) + \langle g(x), y - x\rangle + \tfrac{1}{2}\langle y - x, H(x)(y - x)\rangle + \int_0^1 \int_0^t \langle y - x, [H(x + s(y - x)) - H(x)](y - x)\rangle \, ds\, dt. \tag{1.5.4}$$

Proof. Again considering the univariate functional $\phi(t) := f(x+t(y-x))$, the Fundamental Theorem of Calculus implies

$$\phi(1) = \phi(0) + \phi'(0) + \int_0^1 \phi'(t) - \phi'(0)\, dt \tag{1.5.5}$$

and

$$\phi(1) = \phi(0) + \phi'(0) + \tfrac{1}{2}\phi''(0) + \int_0^1 \int_0^t \phi''(s) - \phi''(0)\, ds\, dt. \tag{1.5.6}$$

Using Proposition 1.4.2 to make the obvious substitutions, (1.5.5) yields (1.5.3) whereas (1.5.6) yields (1.5.4). □

Proposition 1.5.2 provides the means to bound the error in the first and second order approximations of f.

Corollary 1.5.7 *If $x, y \in D_f$ then*

$$|f(y) - f(x) - \langle g(x), y - x\rangle| \leq \|y - x\| \int_0^1 \|g(x + t(y-x)) - g(x)\|\, dt$$

and

$$|f(y) - f(x) - \langle g(x), y - x\rangle - \tfrac{1}{2}\langle y - x, H(x)(y-x)\rangle| \leq \|y - x\|^2 \int_0^1 \int_0^t \|H(x + s(y-x)) - H(x)\|\, ds\, dt.$$

Relying on continuity of g and H, observe that the error in the first-order approximation is $o(\|y - x\|)$ (i.e., tends to zero faster than $\|y - x\|$), whereas the error in the second-order approximation is $o(\|y - x\|^2)$.

Theorem 1.5.1 gives a fundamental theorem of Calculus for a functional f. It will be necessary to have an analogous theorem for g, a theorem which expresses the difference $g(y) - g(x)$ as an integral involving the Hessian. To keep our development coordinate-free, we introduce the following definition:

> The univariate function $t \mapsto v(t) \in \Re^n$, with domain $[a, b]$, is said to be *integrable* if there exists a vector u such that
>
> $$\langle u, w\rangle = \int_a^b \langle v(t), w\rangle\, dt \text{ for all } w \in \Re^n.$$
>
> If it exists, the vector u is uniquely determined (as is not difficult to prove) and is called the *integral* of the function $v(t)$. One uses the notation $\int_a^b v(t)\, dt$ to represent this vector.

Although the definition of the integral is phrased in terms of the inner product $\langle\ ,\ \rangle$, it is independent of the inner product. For if u is the integral as defined by $\langle\ ,\ \rangle$ and if S is pd then for all vectors w,

$$\begin{aligned} \langle u, w\rangle_S &= \langle u, Sw\rangle \\ &= \int_a^b \langle v(t), Sw\rangle \, dt \\ &= \int_a^b \langle v(t), w\rangle_S \, dt. \end{aligned}$$

Following are two useful, elementary propositions.

Proposition 1.5.8 *If the univariate function $t \mapsto v(t) \in \Re^n$, with domain $[a,b]$, is integrable then*

$$\|\int_a^b v(t)\, dt\| \le \int_a^b \|v(t)\|\, dt.$$

Proof. Let $u := \int_a^b v(t)\, dt$. By definition of the integral, for all w we have

$$\langle u, w\rangle = \int_a^b \langle v(t), w\rangle \, dt.$$

In particular, choosing $w = u$ gives

$$\|u\|^2 = \int_a^b \langle v(t), u\rangle \, dt \tag{1.5.9}$$

However,

$$\begin{aligned} \int_a^b \langle v(t), u\rangle \, dt &\le |\int_a^b \langle v(t), u\rangle \, dt| \\ &\le \int_a^b |\langle v(t), u\rangle| \, dt \\ &\le \int_a^b \|v(t)\| \, \|u\| \, dt \\ &= \|u\| \int_a^b \|v(t)\| \, dt. \qquad (1.5.10) \end{aligned}$$

Combining (1.5.9) and (1.5.10) gives

$$\|u\|^2 \le \|u\| \int_a^b \|v(t)\| \, dt.$$

Since $\|u\| = \|\int_a^b v(t)\,dt\|$, the proof is complete. □

Proposition 1.5.11 *If the univariate function $t \mapsto v(t) \in \Re^n$, with domain $[a,b]$, is integrable and if $A : \Re^n \to \Re^m$ is a linear operator then the function $t \mapsto Av(t)$ is integrable and*

$$\int_a^b Av(t)\,dt = A\int_a^b v(t)\,dt.$$

Proof: Observe that for all $w \in \Re^m$ we have

$$\begin{aligned} \langle A\int_a^b v(t)\,dt, w\rangle &= \langle \int_a^b v(t)\,dt, A^* w\rangle \\ &= \int_a^b \langle v(t), A^* w\rangle\,dt \\ &= \int_a^b \langle Av(t), w\rangle\,dt, \end{aligned}$$

where the second equality is by definition of $\int_a^b v(t)\,dt$. □

Next is the fundamental theorem of Calculus for the gradient.

Theorem 1.5.12 *If $x, y \in D_f$ then*

$$g(y) - g(x) = \int_0^1 H(x + t(y-x))(y-x)\,dt.$$

Proof. By definition of the integral, we wish to prove that for all w,

$$\langle g(y) - g(x), w\rangle = \int_0^1 \langle H(x + t(y-x))(y-x), w\rangle\,dt. \tag{1.5.13}$$

Fix arbitrary w and consider the functional

$$\phi(t) := \langle g(x + t(y-x)), w\rangle.$$

The Fundamental Theorem of Calculus asserts

$$\phi(1) - \phi(0) = \int_0^1 \phi'(t)\,dt$$

which, by definition of ϕ, is equivalent to

$$\langle g(y) - g(x), w\rangle = \int_0^1 \phi'(t)\,dt. \tag{1.5.14}$$

Comparing (1.5.14) with (1.5.13), we see that to prove (1.5.13) it suffices to show for arbitrary $0 \leq t \leq 1$ that

$$\phi'(t) = \langle H(u)(y-x), w\rangle \tag{1.5.15}$$

where

$$u := x + t(y-x).$$

Towards proving (1.5.15), recall that $H(u)$ is the unique operator satisfying

$$0 = \lim_{\|\Delta u\|\to 0} \frac{\|g(u+\Delta u) - g(u) - H(u)\Delta u\|}{\|\Delta u\|}. \tag{1.5.16}$$

Thinking of Δu as being $s(y-x)$ where $s \neq 0$, it follows from (1.5.16) that

$$0 = \lim_{s\to 0} \frac{\|g(u+s(y-x)) - g(u) - sH(u)(y-x)\|}{s}. \tag{1.5.17}$$

Since, by Cauchy-Schwarz,

$$\|g(u+s(y-x)) - g(u) - sH(u)(y-x)\|\,\|w\|$$
$$\geq |\langle g(u+s(y-x)) - g(u) - sH(u)(y-x), w\rangle|,$$

(1.5.17) implies

$$0 = \lim_{s\to 0} \frac{\langle g(u+s(y-x)) - g(u) - sH(u)(y-x), w\rangle}{s}.$$

Since

$$\phi(t+s) = \langle g(u+s(y-x)), w\rangle \text{ and } \phi(t) = \langle g(u), w\rangle,$$

we thus have

$$0 = \lim_{s\to 0} \frac{\phi(t+s) - \phi(t) - s\langle H(u)(y-x), w\rangle}{s}$$

from which it is immediate that $\langle H(u)(y-x), w\rangle = \phi'(t)$. Thus, (1.5.15) is established and the proof is complete. □

Proposition 1.5.18 *If $x, y \in D_f$ then*

$$g(y) = g(x) + H(x)(y-x) + \int_0^1 [H(x+t(y-x)) - H(x)](y-x)\,dt.$$

Proof: A simple consequence of Theorem 1.5.12 and

$$\int_0^1 H(x)(y-x)\,dt = H(x)(y-x),$$

an identity which is trivially verified. □

Corollary 1.5.19 *If $x, y \in D_f$ then*

$$\|g(y) - g(x) - H(x)(y-x)\| \leq \|y-x\| \int_0^1 \|H(x+t(y-x)) - H(x)\|\,dt.$$

1.6 Newton's Method

In optimization, Newton's method is an algorithm for minimizing functionals. The idea behind the algorithm is simple: Given a point x in the domain of a functional f, where f is to be minimized, one replaces f by the second-order approximation at x and minimizes the approximation to obtain a new point x_+. One iterates this procedure with x_+ in place of x, and so on, generating a sequence of points which, under certain conditions, converges rapidly to a minimizer of f.

For $x \in D_f$, we denote the second-order – or "quadratic" – approximation at x by

$$q_x(y) := f(x) + \langle g(x), y - x\rangle + \tfrac{1}{2}\langle y - x, H(x)(y - x)\rangle.$$

The domain of q_x is all of $\Re^n$.

Proposition 1.6.1 *The gradient at y of q_x is $g(x) + H(x)(y - x)$ and the Hessian is $H(x)$ (regardless of y).*

Proof: Using the self-adjointness of $H(x)$, is easily established that

$$q_x(y + \Delta y) - q_x(y) - \langle g(x) - H(x)(y - x), \Delta y\rangle = \tfrac{1}{2}\langle \Delta y, H(x)\Delta y\rangle.$$

Proving the gradient is as asserted is thus equivalent to proving

$$\lim_{\|\Delta y\| \to 0} \frac{\langle \Delta y, H(x)\Delta y\rangle}{\|\Delta y\|} = 0,$$

an easily established identity.

Having proven the gradient is as asserted, it is a trivial to prove the Hessian is as asserted. □

Henceforth, assume $H(x)$ is pd. Then q_x is strictly convex and hence is minimized by the point x_+ satisfying $g(x) + H(x)(x_+ - x) = 0$, that is, q_x is minimized by the point

$$x_+ := x - H(x)^{-1}g(x).$$

The "Newton-step at x" is defined to be the difference

$$n(x) := x_+ - x = -H(x)^{-1}g(x).$$

Newton's method steps from x to $x + n(x)$.

We know the second-order approximation is independent of the inner product. Consequently, so is Newton's method. More explicitly, in the inner product $\langle\ ,\ \rangle_S$, the gradient of f at x is $S^{-1}g(x)$, the Hessian is $S^{-1}H(x)$, and so the Newton step is

$$-(S^{-1}H(x))^{-1}S^{-1}g(x) = -H(x)^{-1}g(x).$$

The Newton step is unchanged.

The following theorem is the main tool for analyzing the progress of Newton's method.

Theorem 1.6.2 *If z minimizes f and $H(x)$ is invertible then*

$$\|x_+ - z\| \leq \|x - z\| \|H(x)^{-1}\| \int_0^1 \|H(x + t(z - x)) - H(x)\| \, dt.$$

Proof: Noting $g(z) = 0$, we have

$$\begin{aligned}
\|x_+ - z\| &= \|x - z - H(x)^{-1}g(x)\| \\
&= \|x - z + H(x)^{-1}(g(z) - g(x))\| \\
&= \|x - z + H(x)^{-1} \int_0^1 H(x + t(z - x))(z - x) dt\| \\
&= \|H(x)^{-1} \int_0^1 [H(x + t(z - x)) - H(x)](z - x) \, dt\| \\
&\leq \|x - z\| \|H(x)^{-1}\| \int_0^1 \|H(x + t(z - x)) - H(x)\| \, dt.
\end{aligned}$$

□

Invoking the assumed continuity of the Hessian, the theorem is seen to imply that if $H(z)$ is invertible and x is sufficiently close to z then x_+ will be closer to z than is x.

Now we present a brief discussion of Newton's method and subspaces, as will be important when we consider applications of ipm theory to optimization problems having linear equations among the constraints. Assume L is a subspace of $\Re^n$ and $x \in L \cap D_f$. Let $n|_L(x)$ denote the Newton step for $f|_L$ at x. Since the Hessian of $f|_L$ at x is $P_L H(x)$ and the gradient is $P_L g(x)$, the Newton step $n|_L(x)$ is the vector in L solving

$$P_L H(x) n|_L(x) = -P_L g(x),$$

that is, $n|_L(x)$ is the vector in L for which $H(x)n|_L(x) + g(x)$ is orthogonal to L. In particular, if L is the nullspace of a linear operator $A : \Re^n \to \Re^m$ then $n|_L(x)$ is the vector in $\Re^n$ for which there exists $y \in \Re^m$ satisfying

$$\begin{gathered}
H(x)n|_L(x) + g(x) = A^*y \\
An|_L(x) = 0.
\end{gathered}$$

Computing $n|_L(x)$ (and y) can thus be accomplished by solving a system of $m+n$ equations in $m+n$ variables.

If $H(x)^{-1}$ is readily computed (as it is for functionals f used in ipm's), the size of the system of linear equations to be solved can easily be reduced to m variables. One solves the linear system

$$AH(x)^{-1}A^*y = AH(x)^{-1}g(x)$$

and then computes

$$n|_L(x) = H(x)^{-1}(A^*y - g(x)).$$

In closing this section we remark that of course the error bound given by Theorem 1.6.2 applies to $f|_L$ if z minimizes $f|_L$, $x_+ := x + n|_L(x)$, and the Hessians for f are replaced by Hessians for $f|_L$. In fact, the Hessians need not be replaced by the Hessians for $f|_L$. To verify the replacement need not be done, one notes, for example, that because $H|_L(x) = P_L H(x) P_L$, we have

$$\|H|_L(x)^{-1}\| = 1/\lambda_1' \leq 1/\lambda_1 = \|H(x)^{-1}\|,$$

where λ_1', λ_1, is the smallest eigenvalue of the pd operator $H|_L(x)$, $H(x)$, respectively.

Chapter 2

Basic Interior-Point Method Theory

Throughout this chapter, unless otherwise stated, f refers to a functional having at least the following properties: D_f is open and convex; $f \in \mathcal{C}^2$; $H(x)$ is pd for all $x \in D_f$. In particular, f is strictly convex.

2.1 Intrinsic Inner Products

The functional f gives rise to a family of inner products, $\langle\ ,\ \rangle_{H(x)}$, an inner product for each point $x \in D_f$. These inner products vary continuously with x. In particular, given $\epsilon > 0$, there exists a neighborhood of x consisting of points y with the property that for all vectors $v \neq 0$,

$$1 - \epsilon < \frac{\|v\|_{H(y)}}{\|v\|_{H(x)}} < 1 + \epsilon.$$

We often refer to the inner product $\langle\ ,\ \rangle_{H(x)}$ as the "local inner product (at x)."

In the inner product $\langle\ ,\ \rangle_{H(x)}$, the gradient at y is $H(x)^{-1}g(y)$ and the Hessian is $H(x)^{-1}H(y)$. In particular, the gradient at x is $-n(x)$, the negative of the Newton-step, and the Hessian is I, the identity. Thus, in the local inner product, Newton's method coincides with the "method of steepest descent," i.e., Newton's method coincides with the algorithm which attempts to minimize f by moving in the direction given by the negative of the gradient. (Whereas Newton's method is independent of inner products, the method of steepest descent is not independent because gradients are not independent.)

It appears from our definition that the local inner product potentially depends on the reference inner product $\langle\ ,\ \rangle$ because the Hessian $H(x)$ is w.r.t. that inner product. In fact, the local inner product is independent of the reference inner product. For if the reference inner product is changed to $\langle\ ,\ \rangle_S$, and hence the Hessian is changed to $S^{-1}H(x)$, the resulting local inner product is

$$\langle u, S^{-1}H(x)v\rangle_S = \langle u, SS^{-1}H(x)v\rangle = \langle u, H(x)v\rangle,$$

that is, the local inner product is unchanged.

The independence of the local inner products from the reference inner product shows the local inner products to be intrinsic to the functional f. To highlight the independence of the local products from any reference inner product, we adopt notation which avoids a reference. We denote the local inner product at x by $\langle\ ,\ \rangle_x$. Let $\|\ \|_x$ denote the induced norm. For $y \in D_f$, let $g_x(y)$ denote the gradient at y and let $H_x(y)$ denote the Hessian. Thus, $g_x(x) = -n(x)$ and $H_x(x) = I$. If $A : \Re^n \to \Re^m$ is a linear operator, let A_x^* denotes its adjoint w.r.t. $\langle\ ,\ \rangle_x$. (Of course the adjoint also depends on the inner product on $\Re^m$. That inner product will always be fixed but arbitrary, unlike the intrinsic inner products which vary with x and are not arbitrary, depending on f.)

The reader should be especially aware that we use $g_x(x)$ and $-n(x)$ interchangably, depending on context.

A miniscule amount of the ipm literature is written in terms of the local inner products. Rather, in much of the literature, only a reference inner product is explicit, say, the dot-product. There, the proofs are done by manipulating operators built from Hessians, operators like $H(x)^{-1}H(y)$ and $AH(x)^{-1}A^T$, operators we recognize as being $H_x(y)$ and AA_x^*. An advantage to working in the local inner products is that the underlying geometry becomes evident and, consequently, the operator manipulations in the proofs become less mysterious.

Observe that in the local inner product, the quadratic approximation of f at x is

$$q_x(y) = f(x) - \langle n(x), y - x\rangle_x + \tfrac{1}{2}\|y - x\|_x^2,$$

and its error in approximating $f(y)$ (Corollary 1.5.7) is no worse than

$$\|y - x\|_x^2 \int_0^1 \int_0^t \|I - H_x(x + s(y - x))\|_x \, ds \, dt$$

where the latter norm is the operator norm induced by the local norm. Similarly, the progress made by Newton's method towards approximating a minimizer z (Theorem 1.6.2) is captured by the inequality

$$\|x_+ - z\|_x \leq \|x - z\|_x \int_0^1 \|I - H_x(x + t(z - x))\|_x \, dt.$$

Assume L is a subspace of $\Re^n$ and $x \in L \cap D_f$. Let $P_{L,x}$ denote the operator projecting $\Re^n$ orthogonally onto L, orthogonal w.r.t. $\langle \, , \, \rangle_x$. In the inner product obtained by restricting $\langle \, , \, \rangle_x$ to L, the Hessian of $f|_L$ at x is

$$P_{L,x} H_x(x) = P_{L,x} = I \text{ (the last equality is valid on } L, \text{ not } \Re^n).$$

Consequently, the local inner product on L induced by $f|_L$ is precisely the restriction of $\langle \, , \, \rangle_x$ to L. Letting $g|_{L,x}$ denote the gradient of $f|_L$ w.r.t. the local inner product on L, we thus have

$$n|_L(x) = -g|_{L,x}(x) = -P_{L,x} g_x(x) = P_{L,x} n(x).$$

That is, in the local inner product, the Newton step for $f|_L$ is the orthogonal projection of the Newton step for f.

If L is the nullspace of a surjective linear operator A, the relation

$$n|_L(x) = P_{L,x} n(x) = [I - A_x^*(AA_x^*)^{-1} A] n(x)$$

provides the means to compute $n|_L(x)$ from $n(x)$: One solves the linear system

$$AA_x^* y = -An(x)$$

and then computes

$$n|_L(x) = A_x^* y + n(x).$$

Expressed in terms of an arbitrary inner product $\langle \, , \, \rangle$, the equations become

$$AH(x)^{-1} A^* y = AH(x)^{-1} g(x) \text{ and } n|_L(x) = H(x)^{-1}[A^* y - g(x)],$$

precisely the equations we arrived at in \$1.6 by different reasoning.

2.2 Self-Concordant Functionals

Let $B_x(y, r)$ denote the open ball of radius r centered at y, where radius is measured w.r.t. $\| \; \|_x$. Let $\bar{B}_x(y, r)$ denote the closed ball.

A functional f is said to be *(strongly non-degenerate) self-concordant* if for all $x \in D_f$ we have $B_x(x,1) \subseteq D_f$, and if whenever $y \in B_x(x,1)$ we have

$$1 - \|y - x\|_x \leq \frac{\|v\|_y}{\|v\|_x} \leq \frac{1}{1 - \|y - x\|_x} \quad \text{for all } v \neq 0.$$

Let $\mathcal{SC}$ denote the family of functionals thus defined.

Self-concordant functionals play a central role in the general theory of ipm's, as was first made evident in the pioneering work of Nesterov and Nemirovskii[1]. Although our definition of strongly non-degenerate self-concordant functionals is on the surface quite different from the original definition given in [1], it is in fact equivalent except in assuming $f \in \mathcal{C}^2$ as opposed to the ever-so-slightly stronger assumption in [1] that f is thrice differentiable. The equivalence is shown in §2.5, where it is also shown that our definition can be "relaxed" in a few ways without altering the family of functionals so-defined; for example, the leftmost inequality involving $\|v\|_y/\|v\|_x$ is redundant.

The term "strongly" refers to the requirement $B_x(x,1) \subseteq D_f$. The term "non-degenerate" refers to the Hessians being pd, thereby giving the local inner products. The definition of self-concordant functionals – not necessarily strongly non-degenerate – is a natural relaxation of the above definition, only requiring the Hessians to be psd. However, it is the strongly non-degenerate self-concordant functionals that play the central role in ipm theory and so the relaxation of the definition is best postponed until the reader has in mind a general outline of the theory.

As the parentheses in our definition indicate, for brevity we typically refer to strongly non-degenerate self-concordant functionals simply as "self-concordant functionals."

If a linear functional is added to a self-concordant functional – $x \mapsto \langle c, x\rangle + f(x)$ – the resulting functional is self-concordant because the Hessians are unaffected. Similarly, if one restricts a self-concordant functional f to a subspace L (or to a translation of the subspace), one obtains a self-concordant functional, a simple consequence of the local norms for $f|_L$ being the restrictions of the local norms for f.

The primordial self-concordant barrier functional is the "logarithmic barrier function for the non-negative orthant" having domain $D_f := \Re^n_{++}$ (i.e., the strictly positive orthant). It is defined by $f(x) := -\sum_j \ln x_j$. Since the coordinates of vectors play such a prominent role in the definition of this functional, to prove self-concordancy, it is natural to use the dot product as a reference inner product. Expressing the Hessian $H(x)$ as a matrix, one sees it is diagonal with j^{th} diagonal entry $1/x_j^2$. Consequently, $y \in B_x(x,1)$ is equivalent to

$$\sum_j \left(\frac{y_j - x_j}{x_j}\right)^2 < 1,$$

an inequality which is easily seen to imply $y \in D_f$ as required by the definition of self-concordancy. Moreover, assuming $y \in B_x(x,1)$ and v is an arbitrary vector, we have

$$\begin{aligned} \|v\|_y^2 &= \sum_j \left(\frac{v_j}{y_j}\right)^2 \\ &= \sum_j \left(\frac{v_j}{x_j}\right)^2 \left(\frac{x_j}{y_j}\right)^2 \\ &\leq \|v\|_x^2 \max_j \left(\frac{x_j}{y_j}\right)^2 . \end{aligned}$$

Since

$$\begin{aligned} \frac{x_j}{y_j} &= \left(1 + \frac{y_j - x_j}{x_j}\right)^{-1} \\ &= \sum_{k=0}^{\infty} \left(\frac{x_j - y_j}{x_j}\right)^k \\ &\leq \sum_{k=0}^{\infty} \|y - x\|_x^k \\ &= \frac{1}{1 - \|y - x\|_x}, \end{aligned}$$

the rightmost inequality on $\|v\|_y / \|v\|_x$ in the definition of self-concordancy is proven. The leftmost inequality is proven similarly.

For an LP

$$\begin{array}{ll} \min & c \cdot x \\ \text{s.t.} & Ax = b \\ & x \geq 0, \end{array}$$

the most important self-concordant functionals are those of the form

$$\eta c \cdot x + f|_L(x),$$

where $\eta \geq 0$ is a fixed constant, f is the logarithmic barrier function for the non-negative orthant, and $L := \{x : Ax = b\}$.

Another important self-concordant functional is the "logarithmic barrier function for the cone of pd matrices" in $\mathcal{S}^{n\times n}$. This is the functional defined by $f(X) := -\ln\det(X)$, having domain $\mathcal{S}^{n\times n}_{++}$. To prove self-concordancy, it is natural to rely on the trace product, for which we know $H(X)\Delta X = X^{-1}(\Delta X)X^{-1}$. For arbitrary $Y \in \mathcal{S}^{n\times n}$, keeping in mind that the trace of a matrix depends only on the eigenvalues, we have

$$\begin{aligned}\|Y-X\|_X^2 &= \operatorname{trace}((Y-X)X^{-1}(Y-X)X^{-1})\\ &= \operatorname{trace}(X^{-1/2}(Y-X)X^{-1}(Y-X)X^{-1/2})\\ &= \sum_j (1-\lambda_j)^2,\end{aligned}$$

where $\lambda_1 \leq \ldots \leq \lambda_n$ are the eigenvalues of $X^{-1/2}YX^{-1/2}$. Assuming $\|Y - X\|_X < 1$, all of the values λ_j are thus positive, and hence $X^{-1/2}YX^{-1/2}$ is pd, which is easily seen to be equivalent to Y being pd. Consequently, if $\|Y - X\|_X < 1$ then $Y \in D_f$, as required by the definition of self-concordancy.

Assuming $Y \in B_X(X, 1)$ and letting Q be an orthogonal matrix for which

$$Q^T X^{-1/2} Y X^{-1/2} Q = \Lambda$$

is diagonal, for arbitrary $V \in \mathcal{S}^{n\times n}$ we have

$$\begin{aligned}\|V\|_Y^2 &= \operatorname{trace}(VY^{-1}VY^{-1})\\ &= \operatorname{trace}(X^{-1/2}VY^{-1}VY^{-1}X^{1/2})\\ &= \operatorname{trace}([(X^{-1/2}VX^{-1/2})(X^{1/2}Y^{-1}X^{1/2})]^2)\\ &= \operatorname{trace}(X^{-1/2}VX^{-1/2})(Q\Gamma^{-1}Q^T)]^2)\\ &= \operatorname{trace}([(Q^TX^{-1/2}VX^{-1/2}Q)\Lambda^{-1}]^2)\\ &\leq \tfrac{1}{\lambda_1^2}\operatorname{trace}([Q^TX^{-1/2}VX^{-1/2}Q]^2)\\ &= \tfrac{1}{\lambda_1^2}\operatorname{trace}([X^{-1/2}VX^{-1/2}]^2)\\ &= \tfrac{1}{\lambda_1^2}\operatorname{trace}(VX^{-1}VX^{-1})\\ &= \tfrac{1}{\lambda_1^2}\|V\|_X^2.\end{aligned}$$

Since

$$\begin{aligned}\tfrac{1}{\lambda_1} &= (1+(\lambda_1-1))^{-1}\\ &= \sum_{k=0}^{\infty}(1-\lambda_1)^k\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{k=0}^{\infty}(\sum_{i=1}^{n}(1-\lambda_i)^2)^{k/2} \\
&= \sum_{k=0}^{\infty}\text{trace}([X^{-1/2}(X-Y)X^{-1/2}]^2)^{k/2} \\
&= \sum_{k=0}^{\infty}\text{trace}((X-Y)X^{-1}(X-Y)X^{-1})^{k/2} \\
&= \sum_{k=0}^{\infty}\|Y-X\|_X^k \\
&= \frac{1}{1-\|Y-X\|_X},
\end{aligned}$$

the rightmost inequality on $\|V\|_Y/\|V\|_X$ in the definition of self-concordancy is proven. The leftmost inequality is proven similarly.

For an SDP

$$\begin{array}{ll} \min & C \circ X \\ \text{s.t.} & A(X) = b \\ & X \succeq 0, \end{array}$$

where $A : \mathcal{S}^{n\times n} \to \Re^m$ is a linear operator, the most important self-concordant functionals are those of the form

$$\eta C \circ X + f|_L(X),$$

where $\eta \geq 0$ is a fixed constant, f is the logarithmic barrier function for the cone of pd matrices, and $L := \{X : A(X) = b\}$.

LP can be viewed as a special case of SDP by identifying, in the obvious manner, $\Re^n$ with the subspace in $\mathcal{S}^{n\times n}$ consisting of diagonal matrices. Then the logarithmic barrier function for the pd cone restricts to the logarithmic barrier function for the non-negative orthant. Thus, we were redundant in giving a proof that the logarithmic barrier function for the non-negative orthant is indeed a self-concordant functional. The insight gained from the simplicity of the non-negative orthant justifies the redundancy. In §2.5 we show that the self-conjugacy of each of these two logarithmic barrier functions is a simple consequence of the original definition of self-concordancy due to Nesterov and Nemirovskii[1]. (The original definition is not particularly well-suited for a transparent development of the theory, but it is well-suited for establishing self-conjugacy of logarithmic barrier functions.)

To apply our definition of self-concordancy in developing the theory, it is useful to rephrase it in terms of Hessians. Specifically, since

$$\sup_v \frac{\|v\|_y^2}{\|v\|_x^2} = \sup_v \frac{\langle v, H_x(y)v\rangle_x}{\|v\|_x^2} = \|H_x(y)\|_x$$

and, similarly,

$$\inf_v \frac{\|v\|_y^2}{\|v\|_x^2} = 1/\|H_x(y)^{-1}\|_x,$$

the pair of inequalities in the definition is equivalent to the pair

$$(2.2.1) \qquad \|H_x(y)\|_x, \|H_x(y)^{-1}\|_x \leq \frac{1}{(1 - \|y - x\|_x)^2}.$$

In turn, since

$$\|I - H_x(y)\|_x = \max\{\|H_x(y)\|_x - 1, 1 - 1/\|H_x(y)^{-1}\|_x\}$$

and, similarly,

$$\|I - H_x(y)^{-1}\|_x = \max\{\|H_x(y)^{-1}\|_x - 1, 1 - 1/\|H_x(y)\|_x\},$$

the inequalities (2.2.2) imply

$$(2.2.2) \qquad \|I - H_x(y)\|_x, \|I - H_x(y)^{-1}\|_x \leq \frac{1}{(1 - \|y - x\|_x)^2} - 1.$$

Recalling that $H_x(x) = I$, it is evident from (2.2.2) that to assume self-concordancy is essentially to assume Lipschitz continuity of the Hessians w.r.t. the operator norms induced by the local norms.

An aside for those familar with third differentials: Dividing the quantities on the left and right of (2.2.2) by $\|y - x\|_x$ and taking the limits supremum as y tends to x, suggests when f is thrice differentiable that self-concordancy implies the local norm of the third differential to be bounded by "2." In fact, the converse is also true, that is, a bound of "2" on the local norm of the third differential for all $x \in D_f$, together with the requirement that the local unit balls be contained in the functional domain, imply self-concordancy, as we shall see in §2.5. Indeed, the original definition of self-concordancy in [1] is phrased as a bound on the third differential.

The following proposition and theorem display the simplifying role the conditions of self-concordancy play in analysis. The proposition bounds the error of the quadratic approximation, and the theorem guarantees progress made by Newton's method. Note the elegance of the bounds when compared to the more general Corollary 1.5.7 and Theorem 1.6.2.

Proposition 2.2.3 *If $f \in \mathcal{SC}$, $x \in D_f$ and $y \in B_x(x,1)$ then*

$$|f(y) - q_x(y)| \le \frac{\|y-x\|_x^3}{3(1-\|y-x\|_x)}.$$

Proof: Using Corollary 1.5.7, we have

$$\begin{aligned} |f(y) - q_x(y)| &\le \|y-x\|_x^2 \int_0^1 \int_0^t \|I - H_x(x+s(y-x))\|_x \, ds\, dt \\ &\le \|y-x\|_x^2 \int_0^1 \int_0^t \frac{1}{(1-s\|y-x\|_x)^2} - 1 \, ds\, dt \\ &= \|y-x\|_x^3 \int_0^1 \frac{t^2}{1-t\|y-x\|_x} \, dt \\ &\le \frac{\|y-x\|_x^3}{1-\|y-x\|_x} \int_0^1 t^2 \, dt \\ &= \frac{\|y-x\|_x^3}{3(1-\|y-x\|_x)}. \end{aligned}$$

□

Theorem 2.2.4 *Assume $f \in \mathcal{SC}$ and $x \in D_f$. If z minimizes f and $z \in B_x(x,1)$ then*

$$\|x_+ - z\|_x \le \frac{\|x-z\|_x^2}{1-\|x-z\|_x}.$$

Proof: Using Theorem 1.6.2, simply observe

$$\begin{aligned} \|x_+ - z\|_x &\le \|x-z\|_x \int_0^1 \|I - H_x(x+t(z-x))\|_x \, dt \\ &\le \|x-z\|_x \int_0^1 \frac{1}{(1-t\|y-x\|_x)^2} - 1 \, dt \\ &= \frac{\|x-z\|_x^2}{1-\|x-z\|_x}. \end{aligned}$$

□

The use of the local norm $\| \ \|_x$ in Theorem 2.2.4 to measure the difference $x_+ - z$ makes for a particularly simple proof but does not result in a theorem immediately ready for induction. At x_+, the containment $z \in B_{x_+}(x_+,1)$ is needed to apply the theorem, i.e., a bound on $\|x_+ - z\|_{x_+}$ rather than a bound on $\|x_+ - z\|_x$. Given that the definition of self-concordancy restricts the norms to vary nicely, it is no surprise that the theorem can easily be transformed into a statement ready for induction. For example, substituting into the theorem the inequalities

$$\|x_+ - z\|_z(1 - \|x - z\|_z) \le \|x_+ - z\|_x$$

and

$$\|x - z\|_z(1 - \|x - z\|_z) \le \|x - z\|_x \le \frac{\|x - z\|_z}{1 - \|x - z\|_z},$$

as are immediate from the definition of self-concordancy when $x \in B_z(z, 1)$, we find as a corollary to the theorem that if $\|x - z\|_z < \frac{1}{2}$ then

$$\|x_+ - z\|_z \le \frac{\|x - z\|_z^2}{(1 - \|x - z\|_z)^2(1 - 2\|x - z\|_z)}. \tag{2.2.5}$$

Consequently, if one assumes $\|x - z\|_z < \frac{1}{4}$ then

$$\|x_+ - z\|_z < 4(\|x - z\|_z)^2 \quad (< \tfrac{1}{4}, \text{ so } x_+ \in D_f),$$

and, inductively,

$$\|x_i - z\|_z \le \tfrac{1}{4}(4\|x - z\|_z)^{2^i}, \tag{2.2.6}$$

where $x_1 = x_+, x_2, \ldots$ is the sequence generated by Newton's method. The bound (2.2.6) makes apparent the rapid convergence of Newton's method.

The most elegant proofs of key results in the ipm theory are obtained by phrasing the analysis in terms of $\|n(x)\|_x$ rather than in terms of $\|z - x\|_x$ as was done in Theorem 2.2.4. In this regard, the following theorem is especially useful.

Theorem 2.2.7 *Assume $f \in SC$. If $\|n(x)\|_x < 1$ then*

$$\|n(x_+)\|_{x_+} \le \left(\frac{\|n(x)\|_x}{1 - \|n(x)\|_x}\right)^2.$$

Proof: Assuming $\|n(x)\|_x < 1$, we have

$$\begin{aligned} \|n(x_+)\|_{x_+}^2 &= \|H_x(x_+)^{-1}g_x(x_+)\|_{x_+}^2 \\ &= \langle g_x(x_+), H_x(x_+)^{-1}g_x(x_+)\rangle_x \\ &\le \|H_x(x_+)^{-1}\|_x\|g_x(x_+)\|_x^2. \end{aligned}$$

Since by (2.2.1) we have

$$\|H_x(x_+)^{-1}\|_x \le \frac{1}{(1 - \|n(x)\|_x)^2},$$

we thus have

$$\|n(x_+)\|_{x_+} \le \left(\frac{\|g_x(x_+)\|_x}{1-\|n(x)\|_x}\right)^2.$$

The proof is completed by observing

$$\begin{aligned}\|g_x(x_+)\|_x &= \|\int_0^1 [H_x(x+t\,n(x))-I]n(x)\,dt\|_x \\ &\le \|n(x)\|_x \int_0^1 \|I-H_x(x+t\,n(x))\|_x\,dt \\ &\le \|n(x)\|_x \int_0^1 \frac{1}{(1-t\|n(x)\|_x)^2}-1\,dt \\ &= \frac{\|n(x)\|_x^2}{1-\|n(x)\|_x}.\end{aligned}$$

□

A rather unsatisfying, but unavoidable, aspect of general convergence results for Newton's method is an assumption of x being sufficiently close to a minimizer z, where "sufficiently close" depends explicitly on z. For general functionals, it is impossible to verify that x is indeed sufficiently close to z without knowing z. For self-concordant functionals, we know that the explicit dependence on z of what constitutes "sufficiently close" can take a particularly simple form (e.g., we know x is sufficiently close to z if $\|x-z\|_z < \frac{1}{4}$), albeit a form which appears still to require knowing z. The next proposition provides means to verify proximity to a minimizer without knowing the minimizer.

Proposition 2.2.8 *Assume $f \in \mathcal{SC}$. If $\|n(x)\|_x \le \frac{1}{4}$ then f has a minimizer z and*

$$\|z-x_+\|_x \le \frac{3\|n(x)\|_x^2}{(1-\|n(x)\|_x)^3}.$$

(So, $\|z-x\|_x \le \|n(x)\|_x + 3\|n(x)\|_x^2/(1-\|n(x)\|_x)^3$.)

Proof: We first prove a weaker result, namely, if $\|n(x)\|_x \le \frac{1}{9}$ then f has a minimizer z and $\|x-z\|_x \le 3\|n(x)\|_x$.

Proposition 2.2.3 implies that for all $y \in \bar{B}_x(x, \frac{1}{3})$,

$$|f(y)-q_x(y)| \le \tfrac{1}{6}\|y-x\|_x^2$$

and hence

$$f(y) \ge f(x) - \|n(x)\|_x\|y-x\|_x + \tfrac{1}{3}\|y-x\|_x^2.$$

It follows that if $\|n(x)\|_x \le \frac{1}{9}$ and $\|y-x\|_x = 3\|n(x)\|_x$, then $f(y) \ge f(x)$. However, it is easily proven that whenever a continuous, convex functional f satisfies $f(y) \ge f(x)$ for all y on the boundary of a compact, convex set S and some x in the interior of S, then f has a minimizer in S. Thus, if $\|n(x)\|_x \le \frac{1}{9}$, f has a minimizer z and $\|x-z\|_x \le 3\|n(x)\|_x$.

Now assume $\|n(x)\|_x \le \frac{1}{4}$. Theorem 2.2.7 implies

$$\|n(x_+)\|_{x_+} \le \left(\frac{\|n(x)\|_x}{1-\|n(x)\|_x}\right)^2 \le \tfrac{1}{9}.$$

Applying the conclusion of the preceding paragraph to x_+ rather than to x, we find that f has a minimizer z and $\|z-x_+\|_{x_+} \le 3\|n(x_+)\|_{x_+}$. Thus,

$$\begin{aligned}\|z-x_+\|_x &\le \frac{\|z-x_+\|_{x_+}}{1-\|n(x)\|_x}\\ &\le \frac{3\|n(x_+)\|_{x_+}}{1-\|n(x)\|_x}\\ &\le \frac{3\|n(x)\|_x^2}{(1-\|n(x)\|_x)^3}.\end{aligned}$$

□

We noted that adding a linear functional to a self-concordant functional yields a self-concordant functional. The next two propositions demonstrate other relevant ways for constructing self-concordant functionals.

Proposition 2.2.9 *The set SC is closed under addition, that is, if f_1 and f_2 are self-concordant functionals satisfying $D_{f_1} \cap D_{f_2} \ne \emptyset$ then $f_1 + f_2 : D_{f_1} \cap D_{f_2} \to \Re$ is a self-concordant functional.*

Proof: Let $f := f_1 + f_2$. Assume $x \in D_f$. For all v,

$$\langle v, H(x)v\rangle = \langle v, H_1(x)v\rangle + \langle v, H_2(x)v\rangle,$$

that is

$$\|v\|_x^2 = \|v\|_{x,1}^2 + \|v\|_{x,2}^2.$$

Hence,

$$B_x(x,1) \subseteq B_{x,1}(x,1) \cap B_{x,2}(x,1) \subseteq D_{f_1} \cap D_{f_2} = D_f,$$

as required by the definition of self-concordancy.

Note that whenever a, b, c, d are positive numbers,

$$\min\{\tfrac{a}{c}, \tfrac{b}{d}\} \le \tfrac{a+b}{c+d} \le \max\{\tfrac{a}{c}, \tfrac{b}{d}\}.$$

Consequently, if $y \in D_f$,

$$\min_i \frac{\|v\|_{y,i}^2}{\|v\|_{x,i}^2} \leq \frac{\|v\|_y^2}{\|v\|_x^2} \leq \max_i \frac{\|v\|_{y,i}^2}{\|v\|_{x,i}^2}.$$

Thus, if $\|y - x\|_x < 1$ (and hence $\|y - x\|_{x,i} < 1$),

$$\frac{\|v\|_y}{\|v\|_x} \leq \max_i \frac{1}{1 - \|y - x\|_{x,i}} \leq \frac{1}{1 - \|y - x\|_x},$$

establishing the upper bound on $\|v\|_y / \|v\|_x$ in the defintion of self-concordancy. One establishes the lower bound similarly. □

Proposition 2.2.10 *If $f \in \mathcal{SC}$, $D_f \subseteq \Re^m$ and $A : \Re^n \to \Re^m$ is an injective linear operator then $x \mapsto f(Ax - b)$ is a self-concordant functional, assuming the domain $\{x : Ax - b \in D_f\}$ is non-empty.*

Proof: Denote the functional $x \mapsto f(Ax - b)$ by f'. Assuming $x \in D_{f'}$, one easily verifies from the identity $H'(x) = A^* H(Ax - b)A$ that $H'(x)$ is pd and that $\|v\|_x' = \|Av\|_{Ax-b}$ for all v. In particular,

$$AB_x'(x, 1) - b \subseteq B_{Ax-b}(Ax - b, 1) \subseteq D_f$$

and thus

$$B_x'(x, 1) \subseteq \{y : Ay - b \in D_f\} = D_{f'},$$

as required by the definition of self-concordancy.

If $\|y - x\|_x' < 1$ – hence $\|A(y - x)\|_x < 1$ – and $v \neq 0$ then

$$\begin{aligned} \frac{\|v\|_y'}{\|v\|_x'} &= \frac{\|Av\|_{Ay-b}}{\|Av\|_{Ax-b}} \\ &\leq \frac{1}{1 - \|A(y - x)\|_{Ax-b}} \\ &= \frac{1}{1 - \|y - x\|_x'}, \end{aligned}$$

establishing the upper bound on $\|v\|_y / \|v\|_x$ in the defintion of self-concordancy. One establishes the lower bound similarly. □

Applying Proposition 2.2.10 with the logarithmic barrier function for the non-negative orthant in $\Re^m$, we obtain the self-concordancy of the functional

$$x \mapsto -\sum_i \ln(a_i \cdot x - b_i)$$

whose domain consists of the points x satisfying the strict linear inequality constraints $a_i \cdot x > b_i$. This self-concordant functional is important for LP's with constraints written in the form $Ax \geq b$. It, too, is referred to as a "logarithmic barrier function."

To provide the reader with another (logarithmic barrier) functional with which to apply the above propositions, we mention that $x \mapsto -\ln(1-\|x\|^2)$ is a self-concordant functional with domain the open unit ball. (Verification of self-concordancy is made in §2.5.) Given an ellipsoid $\{x : \|Ax\| < r\}$, it then follows from Proposition 2.2.10 that

$$x \mapsto -\ln(r^2 - \|Ax\|^2)$$

is a self-concordant functional whose domain is the ellipsoid, yet another logarithmic barrier function. For an intersection of ellipsoids, one simply adds the functionals for the individual ellipsoids, as justified by Proposition 2.2.9.

Nesterov and Nemirovskii[1] showed that each open, convex set containing no lines is the domain of a (strongly non-degenerate) self-concordant functional. We give a somewhat different proof of this in §3.7??. Unfortunately, the result is only of theoretical interest. To rely on self-concordant functionals in devising ipm's, one must be able to readily compute their gradients and Hessians. For the self-concordant functionals proven to exist, one cannot say much more than that the gradients and Hessians exist. By contrast, the importance of the various logarithmic barrier functions we have described lies largely in the ease with which their gradients and Hessians can be computed.

Although the values of continuous convex functionals with bounded domains (i.e., bounded w.r.t. any reference norm) are always bounded from below, they need not have minimizers when the domain is open. Such is not the case for self-concordant functionals.

Proposition 2.2.11 *If $f \in SC$ and the values of f are bounded from below then f has a minimizer. (In particular, if D_f is bounded then f has a minimizer.)*

Proof: Assume x satisfies

$$f(x) - \tfrac{1}{108} < \inf_y f(y).$$

Letting $y := x + \frac{1}{3\|n(x)\|_x} n(x)$, Proposition 2.2.3 implies

$$f(y) \leq f(x) - \tfrac{1}{3}\|n(x)\|_x + \tfrac{2}{27}$$

and thus by choice of x,

$$\tfrac{1}{3}\|n(x)\|_x - \tfrac{2}{27} < \tfrac{1}{108},$$

that is,

$$\|n(x)\|_x < 3(\tfrac{2}{27} + \tfrac{1}{108}) = \tfrac{1}{4}.$$

Proposition 2.2.8 then implies f to have a minimizer. □

The conclusion of the next proposition is trivially verified for important self-concordant functionals like those obtained by adding linear functionals to logarithmic barrier functions. Whereas the definition of self-concordancy plays a useful role in both simplifying and unifying the analysis of Newton's method for many functionals important to ipm's, it certainly does not simplify the proof of the property established in the next proposition for those same functionals. Nonetheless, for the theory it is important that the property is possessed by all self-concordant functionals.

Proposition 2.2.12 *Assume $f \in \mathcal{SC}$ and $\tilde{x} \in \partial D_f$, the boundary of D_f. If the sequence $\{x_i\} \subset D_f$ converges to $\tilde{x}$ then* $\liminf_i f(x_i) = \infty$.

Proof: Adding f to a functional $x \mapsto -\ln(\tilde{R} - \|x\|^2)$ where $\tilde{R} > \|\tilde{x}\|$, one obtains a self-concordant functional $\tilde{f}$ for which $D_{\tilde{f}}$ is bounded and for which $\liminf_i \tilde{f}(x_i) = \infty$ iff $\liminf_i f(x_i) = \infty$. Consequently, we may assume D_f is bounded.

Assuming D_f is bounded, we construct from $\{x_i\}$ a sequence $\{y_i\} \subset D_f$ whose limit points lie in ∂D_f and for which

$$f(y_i) \le f(x_i) - \tfrac{1}{60}.$$

Applying the same construction to the sequence $\{y_i\}$, and so on, we will thus conclude that if $\liminf_i f(x_i) < \infty$ then f assumes arbitrarily small values, contradicting the lower boundedness of continuous convex functionals having bounded domains.

Shortly we prove $\liminf_i \|n(x_i)\|_{x_i} \ge \tfrac{1}{5}$. In particular, for sufficiently large i, $y_i := x_i + \frac{1}{5\|n(x_i)\|_{x_i}} n(x_i)$ is well-defined and, from Proposition 2.2.3,

$$f(y_i) \le f(x_i) - \tfrac{1}{25} + \tfrac{1}{50} + \tfrac{1}{300} = f(x_i) - \tfrac{1}{60}.$$

Moreover, all limit points of $\{y_i\}$ lie in ∂D_f. For otherwise, passing to a subsequence of $\{y_i\}$ if necessary, there exists $\epsilon > 0$ such that $B(y_i, \epsilon) \subseteq D_f$ for all i, where the ball is with respect to the reference norm. However, since

y_i and $x_i - (y_i - x_i)$ lie in D_f (because $\|y_i - x_i\|_{x_i} = \frac{1}{5} < 1$), it then follows from convexity of D_f that $B(x_i, \epsilon/2) \subseteq D_f$, contradicting $x_i \to \tilde{x} \in \partial D_f$.

Finally we show $\liminf_i \|n(x_i)\|_{x_i} \geq \frac{1}{5}$. Since D_f is bounded, Proposition 2.2.11 shows f has a minimizer z. Since $B_z(z, 1) \subseteq D_f$ and $x_i \to \tilde{x} \in \partial D_f$, we have $\liminf_i \|x_i - z\|_z \geq 1$. Hence, from the definition of self-concordancy, $\liminf_i \|x_i - z\|_{x_i} \geq \frac{1}{2}$. Because f has a most one minimizer, Proposition 2.2.8 then implies $\liminf_i \|n(x_i)\|_{x_i} \geq \frac{1}{5}$, concluding the proof. □

We close this section with a technical proposition to be called upon later.

If $g(x) + v$ is a vector sufficiently near $g(x)$, there exists $x + u$ close to x such that $g(x + u) = g(x) + v$, a consequence of $H(x)$ being pd and hence invertible. It is useful to quantify "near" and "close" when the inner product is the local inner product, that is, when $H_x(x) = I$ and hence $u \approx v$.

Proposition 2.2.13 *Assume $f \in \mathcal{SC}$ and $x \in D_f$. If $\|v\|_x \leq r$ where $r \leq \frac{1}{4}$, there exists $u \in \bar{B}_x(v, \frac{3r^2}{(1-r)^3})$ such that $g_x(x + u) = g_x(x) + v$.*

Proof: Consider the self-concordant functional

$$y \mapsto -\langle g_x(x) + v, y\rangle_x + f(y), \tag{2.2.14}$$

a functional whose local inner products agree with those of f. Note that a point z' minimizes the functional iff $g_x(z') = g_x(x) + v$. Under the assumption of the proposition, we thus wish to show z' exists and $u := z' - x$ satisfies $u \in \bar{B}_x(v, \frac{3r^2}{(1-r)^3})$.

Since at x, the Newton step for the functional (2.2.14) is v, the assumption $\|v\|_x \leq r$ allows us to apply Proposition 2.2.8, concluding that a minimizer z' does indeed exist and $\|x' - (x + v)\|_x \leq \frac{3r^2}{(1-r)^3}$. □

2.3 Barrier Functionals

A functional f is said to be a *(strongly non-degenerate self-concordant) barrier functional* if $f \in \mathcal{SC}$ and

$$\vartheta_f := \sup_{x \in D_f} \|g_x(x)\|_x^2 < \infty.$$

Let $\mathcal{SCB}$ denote the family of functionals thus defined. We typically refer to elements of $\mathcal{SCB}$ as "barrier functionals."

The definition of barrier functionals is phrased in terms of $\|g_x(x)\|_x$ rather than in terms of the identical quantity $\|n(x)\|_x$ because the importance of barrier functionals for ipm's lies not in applying Newton's method to them directly, but rather, in applying Newton's method to self-concordant functionals built from them. As mentioned before, for an LP

$$\begin{array}{ll} \min & c \cdot x \\ \text{s.t.} & Ax = b \\ & x \geq 0, \end{array}$$

the most important self-concordant functionals are those of the form

$$\eta c \cdot x + f|_L(x), \tag{2.3.1}$$

where $\eta \geq 0$ is a fixed constant, f is the logarithmic barrier function for the non-negative orthant, and $L := \{x : Ax = b\}$.

When they defined barrier functionals, Nesterov and Nemirovskii[1] refered to ϑ_f as "the parameter of the barrier f." Unfortunately, this can be confused with the phrase "barrier parameter" which predates [1] and refers to the constant η in (2.3.1). Consequently, we prefer to call ϑ_f the "complexity value of f," especially because it is the quantity that most often represents f in the complexity analysis of ipm's relying on f.

If one restricts a barrier functional f to a subspace L (or a translation of a subspace), one obtains a barrier functional simply because the local norms for $f|_L$ are the restrictions of the local norms for f and

$$\|g|_{L,x}(x)\|_x = \|P_{L,x} g_x(x)\|_x \leq \|g_x(x)\|_x \leq \sqrt{\varphi_f}.$$

Clearly, $\vartheta_{f|_L} \leq \vartheta_f$.

The primordial barrier functional is the primordial self-concordant functional, i.e., the logarithmic barrier function for the non-negative orthant; $f(x) := -\sum_j \ln x_j$. Relying on the dot product, so that $g(x)$ is the vector with j^{th} entry $1/x_j$ and $H(x)$ is the diagonal matrix with j^{th} diagonal entry $1/x_j^2$, we have

$$\|g_x(x)\|_x^2 = \langle g(x), H(x)^{-1} g(x) \rangle = n.$$

Thus, $\vartheta_f = n$.

Now let f denote the logarithmic barrier function for the cone of pd matrices in $\mathcal{S}^{n \times n}$; $f(X) := -\ln \det(X)$. Relying on the trace product, we have for all $X \in \mathcal{S}^{n\times n}_{++}$,

$$\begin{aligned}\|g_X(X)\|_X^2 &= \langle g(X), H(X)^{-1}g(X)\rangle \\ &= \mathrm{trace}(X^{-1}XX^{-1}X) \\ &= \mathrm{trace}(I) \\ &= n.\end{aligned}$$

Thus, $\vartheta_f = n$.

Finally, let f denote the logarithmic barrier function for the unit ball in $\Re^n$; $f(x) := -\ln(1-\|x\|^2)$ (where $\|\ \| := \langle\ ,\ \rangle^{1/2}$ for some inner product). It is not difficult to verify that for x in the unit ball,

$$g(x) = \frac{2}{1-\|x\|^2}x, \quad H(x)\Delta x = \frac{2}{1-\|x\|^2}\Delta x + \frac{4\langle x, \Delta x\rangle}{(1-\|x\|^2)^2}x,$$

and hence,

$$H(x)^{-1}g(x) = \frac{1-\|x\|^2}{1+\|x\|^2}x.$$

Consequently,

$$\|g_x(x)\|_x^2 = \langle g(x), H(x)^{-1}g(x)\rangle = \frac{2\|x\|^2}{1+\|x\|^2}.$$

It readily follows that $\vartheta_f = 1$, showing the complexity value need not depend on the dimension n.

In §2.2 we noted that if a linear functional is added to a self-concordant functional, the resulting functional is self-concordant because the Hessians are unchanged; the definition of self-concordancy depends on the Hessians alone. By contrast, adding a linear functional to a barrier functional need not result in a barrier functional. For example, consider the the univariate barrier functional $x \mapsto -\ln x$ and the functional $x \mapsto x - \ln x$.

The set $\mathcal{SCB}$, like $\mathcal{SC}$, is closed under addition.

Proposition 2.3.2 *If $f_1, f_2 \in \mathcal{SCB}$ and $D_{f_1} \cap D_{f_2} \neq \emptyset$ then $f := f_1 + f_2 \in \mathcal{SCB}$ (where $D_f = D_{f_1} \cap D_{f_2}$) and $\vartheta_f \le \vartheta_{f_1} + \vartheta_{f_2}$.*

Proof: Assume $x \in D_f$. Let the reference inner product $\langle\ ,\ \rangle$ be the local inner product at x defined by f. Thus, $I = H(x) = H_1(x) + H_2(x)$. In particular, $H_1(x)$ and $H_2(x)$ commute, i.e., $H_1(x)H_2(x) = H_2(x)H_1(x)$. Consequently, so do $H_1(x)^{1/2}$ and $H_2(x)^{1/2}$.

For brevity, let $H_i := H_i(x)$ and $g_i := g_i(x)$ for $i = 1, 2$.

To prove the inequality in the statement of the proposition, it suffices to show

$$\|g_1 + g_2\|^2 \leq \langle g_1, H_1^{-1} g_1\rangle + \langle g_2, H_2^{-1} g_2\rangle.$$

For, by definition, the quantity on the right is bounded by $\vartheta_{f_1} + \vartheta_{f_2}$.

Defining $v_i := H_i^{-1/2} g_i$ for $i = 1, 2$, we have

$$\begin{aligned}
\|g_1 + g_2\|^2 &= \|g_1\|^2 + 2\langle g_1, g_2\rangle + \|g_2\|^2 \\
&= \langle v_1, H_1 v_1\rangle + 2\langle H_1^{1/2} v_1, H_2^{1/2} v_2\rangle + \langle v_2, H_2 v_2\rangle \\
&= \langle v_1, (I - H_2) v_1\rangle + 2\langle H_1^{1/2} v_1, H_2^{1/2} v_2\rangle + \langle v_2, (I - H_1) v_2\rangle \\
&= \langle v_1, v_1\rangle + \langle v_2, v_2\rangle - \|H_2^{1/2} v_1 - H_1^{1/2} v_2\|^2 \\
&\leq \langle v_1, v_1\rangle + \langle v_2, v_2\rangle \\
&= \langle g_1, H_1^{-1} g_1\rangle + \langle g_2, H_2^{-1} g_2\rangle.
\end{aligned}$$

□

The set $\mathcal{SCB}$, like $\mathcal{SC}$, is closed under composition with injective linear maps.

Proposition 2.3.3 *If $f \in \mathcal{SCB}$, $D_f \subseteq \Re^m$ and $A : \Re^n \to \Re^m$ is an injective linear operator then $x \mapsto f(Ax - b)$ is a barrier functional – assuming the domain $\{x : Ax - b \in D_f\}$ is non-empty – and its complexity value does not exceed ϑ_f.*

Proof: Assume $Ax - b \in D_f$. Endow $\Re^n$ with an arbitrary reference inner product and let the reference inner product on $\Re^m$ be the local inner product for f at $Ax - b$. Denoting the functional $x \mapsto f(Ax - b)$ by f', we then have $g'(x) = A^* g(Ax - b)$, $H'(x) = A^* H(Ax - b) A = A^* A$ and $\|g(Ax - b)\|^2 \leq \vartheta_f$. Thus,

$$\begin{aligned}
\langle g'(x), H'(x)^{-1} g'(x)\rangle &= \langle g(Ax - b), A(A^* A)^{-1} A^* g(Ax - b)\rangle \\
&\leq \|A(A^* A)^{-1} A^*\| \, \|g(Ax - b)\|^2 \\
&\leq \vartheta_f.
\end{aligned}$$

(The last inequality is due to the operator being a projection operator; hence the operator has norm equal to one.) The proposition follows. □

With regards to theory, the following proposition is perhaps the most useful tool in establishing properties possessed by all barrier functionals.

Proposition 2.3.4 *Assume $f \in \mathcal{SCB}$. If $x, y \in D_f$ then*

$$\langle g(x), y - x \rangle \leq \vartheta_f.$$

Proof: We wish to prove $\phi'(0) \leq \vartheta_f$ where ϕ is the univariate functional defined by $\phi(t) := f(x + t(y - x))$. In doing so we may assume $\phi'(0) > 0$ and hence, by convexity of ϕ, $\phi'(t) > 0$ for all $t \geq 0$ in the domain of ϕ.

Let $v(t) := x + t(y - x)$. Assuming $t \geq 0$ is in the domain of ϕ,

$$\begin{aligned} \frac{\phi''(t)}{\phi'(t)^2} &= \frac{\langle y - x, y - x \rangle_{v(t)}}{\langle g_{v(t)}(v(t)), y - x \rangle_{v(t)}^2} \\ &\geq \frac{1}{\|g_{v(t)}(v(t))\|_{v(t)}^2} \\ &\geq \frac{1}{\vartheta_f} \end{aligned}$$

and hence for all $s \geq 0$ in the domain of ϕ,

$$\int_0^s \frac{\phi''(t)}{\phi'(t)^2}\, dt \geq \frac{s}{\vartheta_f}.$$

Thus,

$$\left. \frac{-1}{\phi'(t)} \right|_0^s \geq \frac{s}{\vartheta_f},$$

that is,

$$\phi'(s) \geq \frac{\vartheta_f \phi'(0)}{\vartheta_f - s\phi'(0)}.$$

Consequently, $s = \vartheta_f / \phi'(0)$ is not in the domain of ϕ. Since $s = 1$ is in the domain, we have $1 < \vartheta_f / \phi'(0)$. □

The next proposition implies that for each x in the domain of a barrier functional, the ball $B_x(x, 1)$ is, to within a factor of $4\vartheta_f + 1$, the largest among all ellipsoids centered at x which are contained in the domain.

Proposition 2.3.5 *Assume $f \in \mathcal{SCB}$. If $x, y \in D_f$ satisfy $\langle g(x), y - x \rangle \geq 0$ then $y \in B_x(x, 4\vartheta_f + 1)$.*

Proof: Restricting f to the line through x and y, we may assume f is univariate. Viewing the line as $\Re$ with values increasing as one travels from x to y, the assumption $\langle g(x), y - x \rangle \geq 0$ is then equivalent to $g(x) \geq 0$, i.e., $g(x)$ is a non-negative number.

Let v denote the smallest non-negative number for which $\|g_x(x)+v\|_x \geq \frac{1}{4}$. Since $g_x(x) \geq 0$, we have $\|v\|_x \leq \frac{1}{4}$. Applying Proposition 2.2.13, we find there exists u satisfying

$$u \in \bar{B}_x(v, \tfrac{4}{9}) \text{ and } \|g_x(x+u)\|_x = \|g_x(x)+v\|_x \geq \tfrac{1}{4}.$$

Note $\|u\|_x < 1$.

Proposition 2.3.4 implies

$$\begin{aligned} \vartheta_f &\geq \langle g_x(x+u), y-(x+u)\rangle_x \\ &= \langle g_x(x)+v, y-x\rangle_x - \langle g_x(x)+v, u\rangle_x \\ &\geq \tfrac{1}{4}\|y-x\|_x - \langle g_x(x)+v, u\rangle_x, \end{aligned}$$

where the last inequality makes use of $g_x(x)+v$ and $y-x$ both being non-negative. However, since $\|g_x(x)+v\|_x > \frac{1}{4}$ only if $v=0$ (and hence only if $u=0$), we have

$$\langle g_x(x)+v, u\rangle_x \leq \tfrac{1}{4}\|u\|_x < \tfrac{1}{4}.$$

Thus,

$$\vartheta_f > \tfrac{1}{4}\|y-x\|_x - \tfrac{1}{4},$$

from which the proposition is immediate. □

Minimizers of barrier functionals are called *analytic centers*. The following Corollary gives meaning to the term "center."

Corollary 2.3.6 *Assume $f \in \mathcal{SCB}$. If z is the analytic center for f then*

$$B_z(z,1) \subseteq D_f \subseteq B_z(z, 4\vartheta_f + 1).$$

Proof: Since $f \in \mathcal{SC}$, the leftmost containment is by assumption. The rightmost constainment is immediate from Proposition 2.3.5 since $g(z)=0$. □

Corollary 2.3.6 suggests that if one was to choose a single inner product as being especially natural for a barrier functional with bounded domain, the local inner product at the analytic center would be an appropriate choice.

Corollary 2.3.7 *If $f \in \mathcal{SCB}$ then f has an analytic center iff D_f is bounded.*

Proof: Immediate from Proposition 2.2.11 and Corollary 2.3.6. □

In the complexity analysis of ipm's, it is desirable to have barrier functionals with small complexity values. However, there is a positive threshold below which the complexity values of no barrier functionals fall. Nesterov and Nemirovskii[1] prove that $\vartheta_f \geq 1$ for all $f \in \mathcal{SCB}$. To understand why there is indeed a lower bound, assume $\vartheta_f \leq \frac{1}{16}$ for some $f \in \mathcal{SCB}$. Proposition 2.2.8 then implies f has a (unique) minimizer z and all $x \in D_f$ satisfy $\|x - z\|_x \leq \frac{25}{36}$. However, by choosing x so that in the line L through x and z, the distance from x to the boundary of $D_f \cap L$ is smaller than the distance from x to z, the containment $B_x(x, 1) \subseteq D_f$ implies $\|x - z\|_z \geq 1$, a contradiction. Hence $\vartheta_f > \frac{1}{16}$ for all $f \in \mathcal{SCB}$.

Likewise, by Proposition 2.2.8, if $f \in \mathcal{SCB}$ and D_f is unbounded – hence f has no minimizer – then $\|g_x(x)\|_x > \ell := \frac{1}{4}$ for *all* $x \in D_f$. In the unbounded case, Nesterov and Nemirovskii prove the lower bound $\|g_x(x)\|_x \geq \ell := 1$ for all $x \in D_f$.

It is worth noting that a universal lower bound ℓ as in the preceeding paragraph implies a lower bound $n\ell \leq \vartheta_f$ for each barrier functional f whose domain is the non-negative orthant $\Re^n_{++}$. For let e denote the vector of all ones and let e_j denote the j^{th} unit vector. Consider the univariate barrier functional f_j obtained by restricting f to the line through e in the direction e_j. Let g_e denote the gradient of f w.r.t. $\langle\ ,\ \rangle_e$ and let $g_{j,e}$ denote the gradient of f_j w.r.t. the restricted inner product. Since D_{f_j} is unbounded, and hence f_j does not have an analytic center, it is readily proven (without making use of the particular inner product) that $\langle g_{j,e}(e), e_j\rangle_e \leq 0$. Since $g_{j,e}$ and e_j are co-linear (because D_{f_j} is one-dimensional), it follows that

$$\langle g_{j,e}(e), e_j\rangle_e = -\|g_{j,e}(e)\|_e \|e_j\|_e.$$

Noting $\|e_j\|_e \geq 1$ because $e - e_j \notin D_f$, we thus have

$$\langle g_{j,e}(e), e_j\rangle_e \leq -\|g_{j,e}(e)\|_e \leq -\ell.$$

Hence,

$$\begin{aligned} n\ell &\leq \sum_j \langle g_{j,e}(e), -e_j\rangle_e \\ &= \sum_j \langle g_e(e), -e_j\rangle_e \\ &= \langle g_e(e), 0 - e\rangle_e \\ &\leq \vartheta_f, \end{aligned}$$

the last inequality by Proposition 2.3.4.

In light of the two preceding paragraphs, we see that with regards to the complexity value, the logarithmic barrier function for the non-negative orthant $\Re^n_{++}$ is the optimal barrier functional having domain $\Re^n_{++}$. Likewise, viewing $\Re^n$ as a subspace of $\mathcal{S}^{n\times n}$, the logarithmic barrier function for the cone of pd matrices is the optimal barrier functional having that cone as its domain.

For arbitrary inner products, the bounds $\|g_x(x)\|_x \le \sqrt{\vartheta_f}$ imply nothing about the quantities $\|g(x)\|$. However, the bounds do imply bounds on the quantities $\|g_y(x)\|_y$ for all $y \in D_f$. This is the subject of the next proposition. First, a definition.

For x in an arbitrary bounded convex set D, a natural way of measuring the relative nearness of x to the boundary of D, in a manner that is independent of a particular norm, is the quantity known as *the symmetry of D about x*, denoted $\mathrm{sym}(x, D)$. This quantity is defined in terms of the set $\mathcal{L}(x, D)$ consisting of all lines through x which intersect D in an interval of positive length. (If D is lower dimensional, most lines through x will not be in $\mathcal{L}(x, D)$.) If x is an endpoint of $L \cap D$ for some $L \in \mathcal{L}(x, D)$, define $\mathrm{sym}(x, D) := 0$. Otherwise, for each $L \in \mathcal{L}(x, D)$, letting $r(L)$ denote the ratio of the length of the smaller to the larger of the two intervals in $L \cap (D \setminus \{x\})$, define

$$\mathrm{sym}(x, D) := \inf_{L\in\mathcal{L}(x,D)} r(L).$$

Clearly, if D is an ellipsoid centered at x then $\mathrm{sym}(x, D) = 1$, "perfect symmetry." Corollary 2.3.6 implies that if z is the analytic center for a barrier functional f then $\mathrm{sym}(z, D_f) \ge 1/(4\vartheta_f + 1)$.

Proposition 2.3.8 *Assume $f \in \mathcal{SCB}$. If $x \in D_f$ then for all $y \in D_f$,*

$$\|g_y(x)\|_y \le \left(1 + \frac{1}{\mathrm{sym}(x, D_f)}\right)\vartheta_f.$$

Proof: For brevity, let $s := \mathrm{sym}(x, D_f)$. Assuming $x, y \in D_f$, note that $w := y + (1 + s)(x - y) \in \bar{D}_f$, the closure of D_f. Since $B_y(y, 1) \subseteq D_f$ and D_f is convex, we thus have

$$\tfrac{1}{1+s}w + \tfrac{s}{1+s}B_y(y, 1) \subseteq D_f,$$

that is, $B_y(x, \frac{s}{1+s}) \subseteq D_f$. Consequently,

$$\|g_y(x)\|_y \quad = \quad \max_{v\in B_y(x,1)} \langle g_y(x), v - x\rangle_y$$

$$
\begin{aligned}
&= \max_{v \in B_y(x, \frac{s}{1+s})} \tfrac{1+s}{s} \langle g_y(x), v - x \rangle_y \\
&\leq \max_{v \in D_f} \tfrac{1+s}{s} \langle g_y(x), v - x \rangle_y \\
&\leq \tfrac{1+s}{s} \vartheta_f,
\end{aligned}
$$

the last inequality by Proposition 2.3.4. □

At the end of §2.2 we saw that $f(x_i) \to \infty$ if f is a self-concordant functional and $\{x_i\}$ converges to a point in the boundary of D_f. To close this section, we present a proposition that indicates the rate at which $f(x_i)$ goes to ∞ is "slow" if f is a barrier functional.

Proposition 2.3.9 *Assume $f \in \mathcal{SCB}$ and $x \in D_f$. If $y \in \bar{D}_f$ then for all $0 < t \leq 1$,*

$$f(y + t(x - y)) \leq f(x) - \vartheta_f \ln t.$$

Proof: For $s \geq 0$ let $x(s) := y + e^{-s}(x - y)$ and consider the univariate functional $\phi(s) := f(x(s))$. Relying on the chain rule, observe that

$$
\begin{aligned}
\phi(s) &= \phi(0) + \int_0^s \phi'(t)\, dt \\
&= f(x) + \int_0^s \langle g(x(t)), -e^{-s}(x - y) \rangle\, dt \\
&= f(x) + \int_0^s \langle g(x(t)), y - x(t) \rangle\, dt \\
&\leq f(x) + \int_0^s \vartheta_f\, dt \\
&= f(x) + s\vartheta_f,
\end{aligned}
$$

the inequality due to Proposition 2.3.4. Hence, for $0 < t \leq 1$,

$$f(y + t(x - y)) = \phi(-\ln t) \leq f(x) - \vartheta_f \ln(t).$$

□

2.4 Primal Algorithms

The importance of a barrier functional f lies not in itself, but in that it can be used to efficiently solve optimization problems of the form

$$
\text{(2.4.1)} \qquad \begin{array}{ll} \min & \langle c, x \rangle \\ \text{s.t.} & x \in \bar{D}_f, \end{array}
$$

where $\bar{D}_f$ denotes the closure of D_f. Among many other problems, linear programs are of this form. Specifically, restricting the logarithmic barrier function for the non-negative orthant to the space $\{x : Ax = b\}$, we obtain a barrier functional f for which

$$D_f = \{x : Ax = b,\ x \geq 0\}.$$

Similarly for SDP.

Let val denote the optimal value of the optimization problem (2.4.1).

Path-following ipm's solve (2.4.1) by following the *central path*, the path consisting of the minimizers $z(\eta)$ of the self-concordant functionals

$$f_\eta(x) := \eta\langle c, x\rangle + f(x),$$

for $\eta > 0$. It is readily proven when D_f is bounded that the central path begins at the analytic center z of f and consists of the minimizers of the barrier functionals $f|_{L(v)}$ obtained by restricting f to the spaces

$$L(v) := \{x : \langle c, x\rangle = v\}$$

for val $< v < \langle c, z\rangle$. Similarly, when D_f is unbounded, the central path consists of the minimizers of the barrier functionals $f|_{L(v)}$ for val $< v$.

We note the local inner products for the self-concordant functionals f_η are identical with those for f. We observe for each $y \in D_f$, the optimization problem (2.4.1) is equivalent to

$$\begin{array}{ll} \min & \langle c_y, x\rangle_y \\ \text{s.t.} & x \in \bar{D}_f, \end{array}$$

where $c_y := H(y)^{-1}c$. (In other words, w.r.t. $\langle\ ,\ \rangle_y$, the objective vector is c_y.)

The desirability of following the central path is made evident by considering the objective values $\langle c, z(\eta)\rangle$. Since $g(z(\eta)) = -\eta c$, Proposition 2.3.4 implies for all $y \in \bar{D}_f$,

$$\begin{array}{rcl} \langle c, z(\eta)\rangle - \langle c, y\rangle & = & \frac{1}{\eta}\langle g(z(\eta)), y - z(\eta)\rangle \\ & < & \frac{1}{\eta}\vartheta_f \end{array}$$

and hence

$$\langle c, z(\eta)\rangle \leq \text{val} + \tfrac{1}{\eta}\vartheta_f. \tag{2.4.2}$$

Moreover, the point $z(\eta)$ is well-centered in the sense that all feasible points y with objective value at least $\langle c, z(\eta)\rangle$ satisfy $y \in B_{z(\eta)}(z(\eta), 4\vartheta_f + 1)$, a consequence of Proposition 2.3.5.

Path-following ipm's follow the central path approximately, generating points near the central path where "near" is measured by local norms. If a point y is computed for which $\|y - z(\eta)\|_{z(\eta)}$ is small then, relatively, the objective value at y will not be much worse than at $z(\eta)$ and hence (2.4.2) implies a bound on $\langle c, y\rangle$. In fact, if x is an *arbitrary* point in D_f and y is a point for which $\|y - x\|_x$ is small then, relatively, the objective value at y will not be much worse than at x. To make this precise, first observe $B_x(x, 1) \subseteq D_f$ implies $x - tc_x \in D_f$ if $0 \le t < 1/\|c_x\|_x$. Since the objective value at $x - tc_x$ is $\langle c, x\rangle - t\|c_x\|_x^2$ we thus have

$$\|c_x\|_x \le \langle c, x\rangle - \text{val}. \tag{2.4.3}$$

Hence for all $y \in \Re^n$,

$$\begin{aligned} \frac{\langle c, y\rangle - \text{val}}{\langle c, x\rangle - \text{val}} &= 1 + \frac{\langle c_x, y - x\rangle_x}{\langle c, x\rangle - \text{val}} \\ &\le 1 + \frac{\|c_x\|_x \|y - x\|_x}{\langle c, x\rangle - \text{val}} \\ &\le 1 + \|y - x\|_x. \end{aligned}$$

In particular, using (2.4.2),

$$\langle c, y\rangle \le (1 + \|y - x(\eta)\|_{x(\eta)})(\text{val} + \tfrac{1}{\eta}\vartheta_f). \tag{2.4.4}$$

Before discussing algorithms, we record a piece of notation: Let $n_\eta(x)$ denote the Newton step for f_η at x, that is,

$$\begin{aligned} n_\eta(x) &:= -H(x)^{-1}(\eta c + g(x)) \\ &= -(\eta c_x + g_x(x)). \end{aligned}$$

The Barrier Method

"Short-step" ipm's follow the central path most closely, generating sequences of points *all* of which are near the path. We now present and analyze an elementary short-step ipm, the "barrier method."

Assume, initially, we know $\eta_1 > 0$ and x_1 such that x_1 is "near" $z(\eta_1)$, that is, x_1 is near the minimizer for the functional f_{η_1}. In the barrier method, one increases η_1 by a "slight" amount to a value η_2 then applies Newton's method to approximate $z(\eta_2)$, thus obtaining a point x_2. Assuming only one iteration of Newton's method is applied,

$$x_2 := x_1 + n_{\eta_2}(x_1).$$

Continuing this procedure indefinitely (i.e., increasing η, applying Newton's method, increasing η, ...), we have the algorithm known as the barrier method.

One would like η_2 to be much larger than η_1. However, if η_2 is "too" large relative to η_1, Newton's method will fail to approximate $z(\eta_2)$; in fact, it can happen that $x_2 \notin D_f$, bringing the algorithm to a halt. The key ingredient in a complexity analysis of the barrier method is proving that η_2 can be larger than η_1 by a reasonable amount without the algorithm losing sight of the central path.

In analyzing the barrier method, it is most natural to rely on the length of Newton steps to measure proximity to the central path. We will assume x_1 is near $z(\eta_1)$ in the sense that $\|n_{\eta_1}(x_1)\|_{x_1}$ is small. The Newton step taken by the algorithm is $n_{\eta_2}(x_1)$, not $n_{\eta_1}(x_1)$. The relevance of $n_{\eta_1}(x_1)$ for $n_{\eta_2}(x_1)$ is due to the following easily proven relation:

$$n_{\eta_2}(x) = \tfrac{\eta_2}{\eta_1} n_{\eta_1}(x) + (\tfrac{\eta_2}{\eta_1} - 1) g_x(x).$$

In particular,

$$\|n_{\eta_2}(x)\|_x \leq \tfrac{\eta_2}{\eta_1} \|n_{\eta_1}(x)\|_x + |\tfrac{\eta_2}{\eta_1} - 1| \sqrt{\vartheta_f}. \tag{2.4.5}$$

Besides the bound (2.4.5), the other crucial ingredient in the analysis is a bound on $\|n_{\eta_2}(x_2)\|_{x_2}$ in terms of $\|n_{\eta_2}(x_1)\|_{x_1}$. Theorem 2.2.7 provides an appropriate bound: If $\|n_{\eta_2}(x_1)\|_{x_1} < 1$ then

$$\|n_{\eta_2}(x_2)\|_{x_2} \leq \left(\frac{\|n_{\eta_2}(x_1)\|_{x_1}}{1 - \|n_{\eta_2}(x_1)\|_{x_1}} \right)^2 . \tag{2.4.6}$$

Suppose we determine values $\alpha > 0$ and $\beta > 1$ such that if we define

$$\gamma := \alpha\beta + (\beta - 1)\sqrt{\vartheta_f}$$

then $\gamma < 1$ and

$$\left(\frac{\gamma}{1 - \gamma} \right)^2 \leq \alpha.$$

By requiring $\|n_{\eta_1}(x_1)\|_{x_1} \leq \alpha$ and $1 \leq \frac{\eta_2}{\eta_1} \leq \beta$, we then find from (2.4.5) that

$$\|n_{\eta_2}(x_1)\|_{x_1} \leq \gamma,$$

and thus, from (2.4.6),

$$\|n_{\eta_2}(x_2)\|_{x_2} \leq \alpha.$$

Consequently, x_2 will be close to the central path like x_1. Continuing, by requiring $1 \leq \frac{\eta_3}{\eta_2} \leq \beta$, x_3 will be close to the central path, too. And so on. Hence, we will have determined a value β such that if one has an initial point appropriately close to the central path, and if one never increases the barrier parameter from η to more than $\beta\eta$, the barrier method will follow the central path, always generating points close to it.

The reader can verify, for example, that

$$\alpha = \frac{1}{9} \text{ and } \beta := 1 + \frac{1}{8\max\{1, \sqrt{\vartheta_f}\}} \quad \left(= 1 + \frac{1}{8\sqrt{\vartheta_f}}\right)$$

satisfy the relations. Hence we have a "safe" value for β. Relying on it, the algorithm is guaranteed to stay on track. It is a remarkable aspect of ipm's that safe values for quantities like β depend only on the complexity value ϑ_f of the underlying barrier functional f. Concerning LP's,

$$\begin{array}{ll} \min & \langle c, x\rangle \\ \text{s.t.} & Ax = b \\ & x \geq 0, \end{array}$$

if one relies on the logarithmic barrier function for the strictly non-negative orthant $\Re^n_{++}$, then $\gamma = (1 + \frac{1}{8\sqrt{n}})$ is safe, *regardless of A, b and c.*

Assuming at each iteration of the barrier method, the parameter η is never increased by more than $1 + 1/8\sqrt{\vartheta_f}$, we now know that for each x generated by the algorithm, there corresponds $z(\eta)$ which x approximates in that $\|n_\eta(x)\|_x \leq \frac{1}{9}$; hence, by Propositon 2.2.8, $\|x - z(\eta)\|_x \leq \frac{1}{6}$; thus, by the defintion of self-concordancy, $\|x - z(\eta)\|_{z(\eta)} \leq \frac{1}{5}$. All points generated by the algorithm lie within distance $\frac{1}{5}$ of the central path.

Assuming that at each iteration of the barrier method, the parameter η is increased by exactly the factor $1 + 1/8\sqrt{\vartheta_f}$, the number of iterations required to increase the parameter from an initial value η_1 to some value $\eta > \eta_1$ is

$$\begin{aligned} i &= \frac{\ln(\eta/\eta_1)}{\ln(1 + 1/8\sqrt{\vartheta_f})} \\ &\leq 10\sqrt{\vartheta_f}\ln(\eta/\eta_1), \\ &= O(\sqrt{\vartheta_f}\log(\eta/\eta_1)), \end{aligned} \tag{2.4.7}$$

where in the inequality we rely on $\vartheta_f \geq 1$. Hence, from (2.4.4), given $\epsilon > 0$,

$$O\left(\sqrt{\vartheta_f}\log\left(\frac{\vartheta_f}{\epsilon\eta_1}\right)\right) \tag{2.4.8}$$

iterations suffice to produce x satisfying $\langle c, x\rangle \leq \text{val} + \epsilon$.

We have been assuming that an initial point x_1 near the central path is available. What if, instead, we only know some arbitrary point $x' \in D_f$? How might we use the barrier method to solve the optimization problem efficiently? We now describe a simple approach, assuming D_f is bounded and hence f has an analytic center.

Consider the optimization problem obtained by replacing the objective vector c with $-g(x')$. The central path then consists of the minimizers $z'(\nu)$ of the self-concordant functionals

$$f'_\nu(x) := -\nu\langle g(x'), x\rangle + f(x).$$

The point x' is on the central path for this optimization problem. In fact, $x' = z'(\nu)$ for $\nu = 1$.

Let $n'_\nu(x)$ denote the Newton step for f'_ν at x.

Rather than increasing the parameter ν, we decrease it towards zero, following the central path to the analytic center z of f. From there, we switch to following the central path $\{z(\eta)\}$ as before.

We showed η can safely be increased by a factor of $1 + 1/8\sqrt{\vartheta_f}$. Brief consideration of the analysis shows it is also safe to decrease η by a factor $1 - 1/8\sqrt{\vartheta_f}$, and hence, safe to decrease ν by that factor. Thus, to complete our understanding of the difficulty of following the path $\{z'(\nu)\}$, and then the path $\{z(\eta)\}$, it only remains to understand the process of switching paths.

One way to know when it is safe to switch paths is to compute the length of the gradients for f at the points x generated in following the path $\{z'(\nu)\}$. Once one encounters a point x for which, say, $\|g_x(x)\|_x \leq \frac{1}{6}$, one can safely switch paths. For then, by choosing $\eta_1 = 1/12\,\|c_x\|_x$, we find the Newton step for f_{η_1} at x satisfies

$$\|n_{\eta_1}(x)\|_x = \|\eta_1 c_x + g_x(x)\|_x \leq \tfrac{1}{12} + \tfrac{1}{6} = \tfrac{1}{4},$$

and hence, by Proposition 2.2.8, the Newton step takes us from x to a point x_1 for which $\|n_{\eta_1}(x_1)\|_{x_1} \leq \frac{1}{9}$, putting us precisely in the setting of the earlier analysis (where $\alpha = \frac{1}{9}$ was determined safe).

How much will ν have to be decreased from the initial value $\nu = 1$ before we compute a point x for which $\|g_x(x)\|_x \leq \frac{1}{6}$ so that paths can be switched? An answer is found from the relations

$$\begin{aligned} \|g_x(x)\|_x &= \|\nu g_x(x') + n'_\nu(x)\|_x \\ &\leq \nu\|g_x(x')\|_x + \|n'_\nu(x)\|_x \end{aligned}$$

$$\leq \nu\vartheta_f\left(1+\frac{1}{\mathrm{sym}(x',D_f)}\right)+\|n'_\nu(x)\|_x,$$

the last inequality by Proposition 2.3.8. In particular, with $\|n'_\nu(x)\|_x \leq \frac{1}{9}$, ν need only satisfy

$$\nu \leq \frac{1}{18\,\vartheta_f(1+1/\mathrm{sym}(x',D_f))} \tag{2.4.9}$$

in order for $\|g_x(x)\|_x \leq \frac{1}{6}$.

The requirement on ν specified by (2.4.9) gives geometric interpretation to the efficiency of the algorithm in following the path $\{z'(\nu)\}$, beginning with the initial value $\nu = 1$. If the domain D_f is nearly symmetric about the initial point x', not much time will be required to follow the path to a point where we can switch to following the path $\{z(\eta)\}$.

We stipulated that the algorithm switch paths when it encounters x satisfying $\|g_x(x)\|_x \leq \frac{1}{6}$, and we stipulated that one choose the initial value $\eta_1 := 1/12\|c_x\|_x$. Letting

$$V := \sup\{\langle c,x\rangle : x \in D_f\},$$

note 2.4.3 implies

$$\|c_x\|_x \leq V - \mathrm{val}$$

and hence

$$\eta_1 \geq 1/12\,(V-\mathrm{val}).$$

We have now essentially proven the following theorem.

Theorem 2.4.10 *Assume $f \in \mathcal{SCB}$ and D_f is bounded. Assume $x' \in D_f$, a point at which to initiate the barrier method. If $0 < \epsilon < 1$, then within*

$$O\left(\sqrt{\vartheta_f}\log\left(\frac{\vartheta_f}{\epsilon\,\mathrm{sym}(x',D_f)}\right)\right)$$

iterations of the method, all points x computed thereafter satisfy

$$\frac{\langle c,x\rangle - \mathrm{val}}{V-\mathrm{val}} \leq \epsilon.$$

Consider the following modification to the algorithm: Choose $V' > \langle c,x'\rangle$. Rather than relying on f, rely on the barrier functional

$$x \mapsto f(x) - \ln(V' - \langle c,x\rangle),$$

a functional whose domain is

$$D_f \cap \{x : \langle c, x\rangle < V'\} \qquad (2.4.11)$$

and whose complexity value does not exceed $\vartheta_f + 1$. In the theorem, the quantity $V -$ val is then replaced by the potentially much smaller quantity $V' -$ val. Of course the quantity $\text{sym}(x', D_f)$ must then be replaced by the symmetry of the set (2.4.11) about x'.

Finally, we highlight an implicit assumption underlying our analysis, namely, the complexity value ϑ_f is known. The value is used to safely increase the parameter η. What is actually required is an upper bound $\vartheta \geq \vartheta_f$. If one relies on an upper bound ϑ rather than the precise complexity value ϑ_f then ϑ_f in the theorem must be replaced by ϑ.

Except for ϑ_f, none of the quantities appearing in the theorem are assumed to be known or approximated. The quantities appear naturally in the analysis of the algorithm but the algorithm itself does not rely on the quantities.

No ipm's have proven complexity bounds which are better than (2.4.8), even in the restricted setting of linear programming. Nonetheless, the barrier method is not considered to be practically efficient relative to some other ipm's, especially relative to primal-dual methods (discussed in Chapter 3). The barrier method is an excellent algorithm with which to begin one's understanding of ipm's, and it is often the perfect choice for concise complexity theory proofs, but it is not one of the ipm's that appear in widely used software.

The Long-Step Barrier Method

One of the barrier method's shortcomings is obvious, being implicit in the terminology "short-step algorithm." Although it is *always* safe to increase η by a factor $1 + 1/8\sqrt{\vartheta_f}$ with each iteration, that increase is small if ϑ_f is large. No doubt, for many instances, a much larger increase is safe.

There is a trivial manner in which to modify the barrier method in hopes of having a more practical algorithm. Rather than increase η by the safe amount, increase it by much more, apply (perhaps several iterations of) Newton's method, and check (say, using Proposition 2.2.8) if the computed point is near the desired minimizer. If not, increase η by a smaller amount and try again.

A more interesting and more practical modification of the barrier method is known as the "long-step barrier method." In this version, one increases η by an arbitrarily large amount but does not take Newton steps. Instead, the Newton steps are used as directions for "exact line searches," as we now describe.

Assume as before that we have an initial value $\eta_1 > 0$ and a point x_1 approximating $z(\eta_1)$. Choose η_2 larger than η_1, perhaps significantly larger. In search of a point x_2 which approximates $z(\eta_2)$, the algorithm will generate a finite sequence of points

$$y_1 := x_1, y_2, \ldots, y_{K-1}, y_K,$$

then let $x_2 := y_K$. At each point y_k, the algorithm will determine if the point is close to $z(\eta_2)$ by, say, checking whether $\|n_{\eta_2}(y_k)\|_{y_k} \leq \frac{1}{4}$. (We choose the specific value $\frac{1}{4}$ because it is the largest value for which Proposition 2.2.8 applies.) The point y_K will be the first point that is determined to satisfy this inequality.

To compute y_{k+1} from y_k, the algorithm minimizes the univariate functional

$$t \mapsto f(y_k + t\, n_{\eta_2}(y_k)). \tag{2.4.12}$$

This is the step in the algorithm to which the phrase "exact line search" alludes. "Line" refers to the functional being univariate. "Exact" refers to an assumption that the exact minimizer is computed, certainly an exageration, but an assumption needed to keep the complexity analysis succinct. Letting t_{k+1} denote the exact minimizer, define

$$y_{k+1} := y_k + t_{k+1} n_{\eta_2}(y_k),$$

thus ending our description of the long-step barrier method.

The short-step barrier method is confined to making slow but sure progress. The long-step method is more adventurous, having the potential for much quicker progress.

Clearly, the complexity analysis of the long-step barrier method revolves around determining an upper bound on K in terms of the ratio η_2/η_1. We now undertake the task of determining such a bound.

We begin by determining an upper bound on the difference

$$\rho := f_{\eta_2}(x_1) - f_{\eta_2}(z(\eta_2)).$$

Then we show that $f_{\eta_2}(y_k) - f_{\eta_2}(y_{k+1})$ is bounded below by a positive amount τ independent of k, that is, each exact line search decreases the value of f_{η_2} by at least a certain amount. Consequently $K \leq \rho/\tau$. Proofs like this – showing a certain functional decreases by at least a certain amount with each iteration – are common in the ipm literature.

In proving an upper bound on the difference ρ, we make use of the fact that for any convex functional f and $x, y \in D_f$, one has

$$f(x) - f(y) \le -\langle g(x), y - x \rangle. \tag{2.4.13}$$

The upper bound on ρ is obtained by adding upper bounds for

$$\rho_1 := f_{\eta_2}(x_1) - f_{\eta_2}(z(\eta_1)) \text{ and } \rho_2 := f_{\eta_2}(z(\eta_1)) - f_{\eta_2}(z(\eta_2)).$$

Assuming x_1 is close to $z(\eta_1)$ in the sense that $\|n_{\eta_1}(x_1)\|_{x_1} \le \frac{1}{4}$, Proposition 2.2.8 implies $\|x_1 - z(\eta_1)\|_{x_1} < \frac{1}{2}$. Thus, applying (2.4.13) to the functional f_{η_2},

$$\begin{aligned} \rho_1 &\le \langle n_{\eta_2}(x_1), z(\eta_1) - x_1 \rangle_{x_1} \\ &= \tfrac{\eta_2}{\eta_1} \langle n_{\eta_1}(x_1), z(\eta_1) - x_1 \rangle_{x_1} \\ &\quad + (\tfrac{\eta_2}{\eta_1} - 1) \langle g_{x_1}(x_1), z(\eta_1) - x_1 \rangle_{x_1} \\ &\le \tfrac{\eta_2}{\eta_1} \tfrac{1}{4} \tfrac{1}{2} + (\tfrac{\eta_2}{\eta_1} - 1) \tfrac{1}{2} \sqrt{\vartheta_f} \\ &\le \tfrac{\eta_2}{\eta_1} \sqrt{\vartheta_f}. \end{aligned}$$

Similarly, for all $y \in D_f$,

$$\begin{aligned} f_{\eta_2}(z(\eta_1)) - f_{\eta_2}(y) &\le \langle n_{\eta_2}(z(\eta_1)), y - z(\eta_1) \rangle_{z(\eta_1)} \\ &= \tfrac{\eta_2}{\eta_1} \langle n_{\eta_1}(z(\eta_1)), y - z(\eta_1) \rangle_{z(\eta_1)} \\ &\quad + (\tfrac{\eta_2}{\eta_1} - 1) \langle g_{z(\eta_1)}(z(\eta_1)), y - z(\eta_1) \rangle_{z(\eta_1)} \\ &= (\tfrac{\eta_2}{\eta_1} - 1) \langle g_{z(\eta_1)}(z(\eta_1)), y - z(\eta_1) \rangle_{z(\eta_1)}, \end{aligned}$$

the final equality because $z(\eta_1)$ minimizes f_{η_1} and hence $n_{\eta_1}(z(\eta_1)) = 0$. Thus,

$$\begin{aligned} \rho_2 &\le (\tfrac{\eta_2}{\eta_1} - 1) \langle g(z(\eta_1)), z(\eta_2) - z(\eta_1) \rangle_{z(\eta_1)} \\ &\le \tfrac{\eta_2}{\eta_1} \vartheta_f, \end{aligned}$$

the last inequality by Proposition 2.3.4.

Now we show $f_{\eta_2}(y_k) - f_{\eta_2}(y_{k+1})$ is bounded below by a positive amount τ independent of k.

If the algorithm proceeds from y_k to y_{k+1}, it is because y_k happens not to be appropriately close to $z(\eta_2)$, i.e., it happens that $\|n_{\eta_2}(y_k)\|_{y_k} > \frac{1}{4}$.

Letting $\tilde{t} := 1/5\|n_{\eta_2}(y_k)\|_{y_k}$ and $\tilde{y} := y_k + \tilde{t}n_{\eta_2}(y_k)$, Proposition 2.2.3 then implies

$$\begin{aligned} f_{\eta_2}(\tilde{y}) &\leq f_{\eta_2}(y_k) - \tfrac{1}{4}\tfrac{1}{5} + \tfrac{1}{2}(\tfrac{1}{5})^2 + \tfrac{(1/5)^3}{3(1-1/5)} \\ &\leq f_{\eta_2}(y_k) - \tfrac{1}{40}. \end{aligned}$$

Since t_{k+1} minimizes the functional (2.4.12), we thus have

$$f_{\eta_2}(y_k) - f_{\eta_2}(y_{k+1}) \geq \tau := \tfrac{1}{40}.$$

Finally,

$$K \leq \frac{\rho}{\tau} \leq \frac{\rho_1 + \rho_2}{\tau} \leq \tfrac{40\,\eta_2}{\eta_1}(\vartheta_f + \sqrt{\vartheta_f}).$$

It follows that if one fixes a positive constant $\kappa > 1$ and always chooses successive values η_i, η_{i+1} to satisfy $\eta_{i+1} = \kappa\eta_i$, the number of points generated by the long-step barrier method (i.e., the number of exact line searches) in increasing the parameter from an initial value η_1 to some value $\eta > \eta_1$ is

$$O(\kappa\vartheta_f \log(\eta/\eta_1)).$$

No better bound is known for the long-step method. Fixing κ (say, $\kappa = 100$), we obtain the bound

$$O(\vartheta_f \log(\eta/\eta_1)).$$

This bound is worse than the analogous bound (2.4.7) for the short-step method by a factor $\sqrt{\vartheta_f}$. It is one of the ironies of t he ipm literature that algorithms which are more efficient in practice often have slightly worse complexity bounds.

A Predictor-Corrector Method

The Newton step $n_\eta(x) := -\eta c_x - g_x(x)$ for the barrier method can be viewed as the sum of two steps, one of which predicts the tangential direction of the central path and the other of which corrects for the discrepancy between the tangential direction and the actual position of the (curving) path.

The corrector step is the Newton step at x for the barrier functional $f|_{L(v)}$ where $v = \langle c, x\rangle$ and

$$L(v) := \{y : \langle c, y\rangle = v\}.$$

Thus, the correcting step is $n|_{L(v)}(x)$, this being the orthogonal projection of the Newton step $n(x)$ for f onto the subspace $L(0)$ (where "orthogonal" is w.r.t. $\langle\ ,\ \rangle_x$). In the literature, the corrector step is often referred to as the "centering direction." It aims to move from x towards the point on the central path having the same objective value as x.

Since the multiples of c_x $(= H(x)^{-1}c)$ form the orthogonal complement of $L(0)$, the difference $n(x) - n|_{L(v)}(x)$ is a multiple of c_x and hence so is the predictor step

$$n_\eta(x) - n|_{L(v)}(x) = -\eta c_x + (n(x) - n|_{L(v)}(x)).$$

The vector $-c_x$ predicts the tangential direction of the central path near x. If x is on the central path, the vector $-c_x$ is exactly tangential to the path, pointing in the direction of decreasing objective values. In the literature, $-c_x$ is often referred to as the "affine-scaling direction." With regards to $\langle\ ,\ \rangle_x$, it is the direction in which one would move to decrease the objective value most quickly.

Whereas the barrier method combines a predictor step and a corrector step in one step, predictor-corrector methods separate the two types of steps. After a predictor step, several corrector steps might be applied. In practice, predictor-corrector methods tend to be substantially more efficient than the barrier method, but the (worst-case) complexity bounds that have been proven for them are worse.

Perhaps the most natural predictor-corrector method is based on moving in the predictor direction a fixed fraction of the distance towards the boundary and then re-centering via exact line searches. We now formalize and analyze such an algorithm.

Fix σ satisfying $0 < \sigma < 1$. Assume x_1 is near the central path. Let $v_1 := \langle c, x_1\rangle$. The algorithm is assumed to first compute

$$s_1 := \sup\{s : x - sc_{x_1} \in D_f\}$$

and then let $y_1 := x - \sigma s_1 c_{x_1}$. Thus, $-\sigma s_1 c_{x_1}$ is the predictor step. Let

$$v_2 := \langle c, y_1\rangle = v_1 - \sigma s_1 \|c_{x_1}\|_{x_1}^2.$$

Beginning with y_1, the algorithm takes corrector steps, moving towards the point z_2 on the central path with objective value v_2 by using the Newton steps for the functional $f|_{L(v_2)}$ as directions in performing exact line searches. Precisely, given y_k, the algorithm is assumed to compute exactly the minimizer t_{k+1} for the univariate functional

$$t \mapsto f(y_k + tn|_{L(v_2)}(y_k)),$$

and then let

$$y_{k+1} := y_k + t_{k+1} n|_{L(v_2)}(y_k).$$

When the first point y_K is encountered for which $\|n|_{L(v_2)}(y_K)\|_{y_K}$ is appropriately small, the algorithm lets $x_2 := y_K$ and takes a predictor step from x_2, relying on the same value σ as in the predictor step from x_1. The predictor step is followed by corrector steps, and so on.

Typically, σ is chosen very nearly equal to 1, say $\sigma = .99$. In the following analysis, we assume $\sigma \geq \frac{1}{4}$.

In analyzing the predictor-corrector method, we determine an upper bound on the number K of exact line searches made in moving from x_1 to x_2, and we determine a lower bound on progress made in decreasing the objective value by moving from x_1 to x_2.

For the analysis, we consider $\|n|_{L(v)}(x)\|_x \leq \frac{1}{13}$ to be the criterion for claiming x to be close to the point z on the central path satisfying $\langle c, z\rangle = \langle c, x\rangle$. The specific value $\frac{1}{13}$ is chosen so we will be in position to rely on the barrier method analysis. To see how it puts us in position to rely on that analysis, let z_1 denote the minimizer of $f|_{L(v_1)}$ and let $\eta_1 > 0$ denote the value for which $z(\eta_1) = z_1$. We claim that $\|n|_{L(v_1)}(x_1)\|_{x_1} \leq \frac{1}{13}$ implies $\|n_{\eta_1}(x_1)\|_{x_1} \leq \frac{1}{9}$, precisely the criteria we assumed x_1 to satisfy in the barrier method. For if $\|n|_{L(v_1)}(x_1)\|_{x_1} \leq \frac{1}{13}$ then Proposition 2.2.8 applied to $f|_{L(v_1)}$ implies

$$\|z_1 - x_1\|_{x_1} \leq \tfrac{1}{10}. \tag{2.4.14}$$

Since $z_1 = z(\eta_1)$, applying Theorem 2.2.4 to f_{η_1} then yields $\|n_{\eta_1}(x_1)\|_{x_1} \leq \frac{1}{9}$.

The barrier method moves from x_1 to $x_1 + n_{\eta_2}(x_1)$ where $\eta_2 = (1 + 1/8\sqrt{\vartheta_f})\eta_2$. (It is not necessarily the case that $z_2 = z(\eta_2)$ where z_2 minimizes $f|_{L(v_2)}$.) The length of the barrier method step is thus

$$\begin{aligned}\|n_{\eta_2}(x_1)\|_{x_1} &= \|\tfrac{\eta_2}{\eta_1} n_{\eta_1}(x_1) + (\tfrac{\eta_2}{\eta_1} - 1) g_{x_1}(x_1)\|_{x_1} \\ &\leq \tfrac{9}{8}\tfrac{1}{9} + \tfrac{1}{8} \\ &= \tfrac{1}{4}.\end{aligned}$$

Consequently, in one step of the barrier method, the objective value decreases by at most $\frac{1}{4}\|c_x\|_x$.

Assuming $\sigma \geq \frac{1}{4}$, in the predictor-corrector method the predictor step is in the direction $-c_{x_1}$ and has length at least $\frac{1}{4}$, a consequence of $B_{x_1}(x_1, 1) \subseteq D_f$. Thus, $v_1 - v_2 \geq \frac{1}{4}\|c_{x_1}\|_{x_1}$. Hence, in moving from x_1 to x_2, the progress made by the predictor-corrector method in decreasing the objective value is at least as great as the progress made by the barrier method in taking one step from x_1. Of course in moving from x_1 to x_2, the predictor-corrector method might require several exact line searches. We now bound the number K of exact line searches.

Analogous to our analysis for the long-step barrier method, we obtain an upper bound on K by dividing an upper bound on $f(y_1) - f(z_2)$ by a lower bound on the differences $f(y_k) - f(y_{k+1})$.

The lower bound on the differences $f(y_k) - f(y_{k+1})$ is proven exactly as was the lower bound for the differences $f_{\eta_2}(y_k) - f_{\eta_2}(y_{k+1})$ in our analysis of the long-step barrier method: Assuming $\|n|_{L(v_2)}(y_k)\|_{y_k} > \frac{1}{13}$ (as is the case if the algorithm proceeds to compute y_{k+1}), one relies on Propositon 2.2.3, now applied to the functional $f|_{L(v_2)}$, to show $f(y_k) - f(y_{k+1})$ is bounded below independently of k.

To obtain an upper bound on $f(y_1) - f(z_2)$, one can use the relation

$$f(y_1) - f(z_2) = (f(y_1) - f(x_1)) + (f(x_1) - f(z_1)) + (f(z_1) - f(z_2)).$$

Proposition 2.3.9 and the definition of y_1 imply

$$f(y_1) - f(x_1) \leq -\vartheta_f \ln(1 - \sigma).$$

Relation (2.4.13) applied to $f|_{L(v_1)}$, together with (2.4.14), give

$$f(x_1) - f(z_1) \leq \|n|_{L(v_1)}(x_1)\|_{x_1} \|z_1 - x_1\|_{x_1} \leq \tfrac{1}{13}\tfrac{1}{10}.$$

Finally, (2.4.14) applied to f gives

$$f(z_1) - f(z_2) \leq \langle \eta_1 c, z_2 - z_1 \rangle \leq 0.$$

In all,

$$f(z_1) - f(z_2) \leq \vartheta_f \ln(\tfrac{1}{1-\sigma}) + \tfrac{1}{130}.$$

Combined with the constant lower bound on the differences $f(y_k) - f(y_{k+1})$ we thus find the number K of exact line searches performed in moving from x_1 to x_2 satisfies

$$K = O\left(\vartheta_f \log\left(\frac{1}{1-\sigma}\right)\right).$$

Having shown the progress made by the predictor-corrector method in decreasing the objective value is at least as great as the progress made by the barrier method in taking one step from x_1, we obtain complexity bounds for the predictor-corrector method which are greater than the bounds for the barrier method by a factor K, that is, by a factor ϑ_f (assuming σ fixed; say, $\sigma = .99$). The bounds are greater than the bounds for the long-step barrier method by a factor $\sqrt{\vartheta_f}$. No better bounds are known for the predictor-corrector method.

2.5 Matters of Definition

There are various equivalent ways to define self-concordant functionals. Our definition is geometric and simple to employ in theory, but it is not the original definition due to Nesterov and Nemirovskii[1]. In this section we consider various equivalent definitions of self-concordancy, including the original definition. We close the section with a brief discussion of the term "strongly non-degenerate," a term we have supressed thus far.

Unless otherwise stated, we assume that $f \in C^2$, D_f is open and convex, and $H(x)$ is pd for all $x \in D_f$.

For ease of reference, we recall our definition of self-concordancy:

> A functional f is said to be *(strongly non-degenerate) self-concordant* if for all $x \in D_f$ we have $B_x(x,1) \subseteq D_f$, and if whenever $y \in B_x(x,1)$ we have
> $$1 - \|y-x\|_x \le \frac{\|v\|_y}{\|v\|_x} \le \frac{1}{1-\|y-x\|_x} \quad \text{for all } v \neq 0.$$

Recall $\mathcal{SC}$ denotes the family of functionals thus defined.

An important property is establishing the equivalence of various definitions of self-concordancy is the "transitivity" of the condition

$$\forall v \neq 0, \quad \frac{\|v\|_y}{\|v\|_x} \le \frac{1}{1-\|y-x\|_x}. \tag{2.5.1}$$

Specifically, if x, y and z are co-linear with y between x and z, if x and y satisfy (2.5.1), if y and z satisfy the analogous inqualities

$$\forall v \neq 0, \quad \frac{\|v\|_z}{\|v\|_y} \le \frac{1}{1-\|z-y\|_y}, \tag{2.5.2}$$

and if $\|z-x\|_x < 1$, then

$$\forall v \neq 0, \quad \frac{\|v\|_z}{\|v\|_x} \le \frac{1}{1-\|z-x\|_x}. \tag{2.5.3}$$

To establish this transitivity, note (2.5.1) implies

$$\|z-y\|_y \le \frac{\|z-y\|_x}{1-\|y-x\|_x}.$$

Substituting into (2.5.2) gives for all $v \neq 0$,

$$\begin{aligned} \frac{\|v\|_z}{\|v\|_y} &\le \frac{1-\|y-x\|_x}{1-\|y-x\|_x-\|z-y\|_x} \\ &= \frac{1-\|y-x\|_x}{1-\|z-x\|_x}, \end{aligned} \tag{2.5.4}$$

the equality relying on the co-linearity of x, y and z. Note (2.5.3) is immediate from (2.5.1) and (2.5.4). Hence the transitivity.

Our first modification to the definition of self-concordancy is to show the redundancy of the leftmost inequality in the definition.

Proposition 2.5.5 *Assume f is such that for all $x \in D_f$ we have $B_x(x,1) \subseteq D_f$, and is such that whenever $y \in B_x(x,1)$ we have*

$$\frac{\|v\|_y}{\|v\|_x} \leq \frac{1}{1 - \|y-x\|_x} \quad \textit{for all } v \neq 0. \tag{2.5.6}$$

Then

$$1 - \|y-x\|_x \leq \frac{\|v\|_y}{\|v\|_x} \quad \textit{for all } v \neq 0.$$

Proof: Since

$$\|H_x(y)^{-1}\|_x = \sup_{v \neq 0} \frac{\|v\|_x^2}{\|v\|_y^2},$$

it suffices to show for all x, y and $0 < \epsilon < 1$,

$$\|y-x\|_x \leq \frac{1}{1+\epsilon} \quad \Rightarrow \quad \|H_x(y)^{-1}\|_x \leq \frac{1}{(1-(1+\epsilon)\|y-x\|_x)^2}. \tag{2.5.7}$$

Towards establishing the implication (2.5.7), note that from (2.5.6), whenever points y and z satisfy $\|z-y\|_y \leq \frac{\epsilon}{1+\epsilon}$ we have

$$\|z-y\|_z \leq \frac{\|z-y\|_y}{1-\|z-y\|_y} \leq (1+\epsilon)\|z-y\|_y. \tag{2.5.8}$$

Letting $\lambda_{\min}$ denote the minimum eigenvalue of $H_y(z)$, also note that since $H_y(z)$ is self-adjoint w.r.t. both $\langle\ ,\ \rangle_y$ and $\langle\ ,\ \rangle_z$, we have

$$\|H_y(z)^{-1}\|_y = 1/\lambda_{\min} = \|H_y(z)^{-1}\|_z = \|H_z(y)\|_z.$$

Recalling

$$\|H_z(y)\|_z = \sup_{v \neq 0} \frac{\|v\|_y^2}{\|v\|_z^2},$$

we thus have by (2.5.6) and (2.5.8),

$$\|H_y(z)^{-1}\|_y \leq \frac{1}{(1-\|y-z\|_z)^2} \leq \frac{1}{(1-(1+\epsilon)\|y-z\|_z)^2}.$$

In other words,

$$(2.5.9)\quad \|z-y\|_y \le \frac{\epsilon}{1+\epsilon} \quad \Rightarrow \quad \|H_y(z)^{-1}\|_y \le \frac{1}{(1-(1+\epsilon)\|z-y\|_y)^2},$$

which would give the desired implication (2.5.7) except that $\frac{\epsilon}{1+\epsilon}$ replaces $\frac{1}{1+\epsilon}$.

In light of (2.5.9), to prove (2.5.7) it suffices to establish that for all y ($\neq x$) satisfying

$$(2.5.10)\quad \|y-x\|_x < \frac{1}{1+\epsilon} \text{ and } \|H_x(y)^{-1}\|_x \le \frac{1}{1-(1+\epsilon)\|y-x\|_x)^2},$$

if we define $z := y + t(y-x)$ where t is chosen so that

$$\|z-y\|_x = \epsilon(\tfrac{1}{1+\epsilon} - \|y-x\|_x),$$

then

$$\|H_x(z)^{-1}\|_x \le \frac{1}{(1-(1+\epsilon)\|z-x\|_z)^2}.$$

Assuming x, y and z satisfy the assumptions just stated, using (2.5.6) we have

$$\begin{aligned}
\|z-y\|_y &\le \frac{\|z-y\|_x}{1-\|y-x\|_x} \\
&= \frac{\frac{\epsilon}{1+\epsilon}[1-(1+\epsilon)\|y-x\|_x}{1-\|y-x\|_x} \\
&\le \frac{\epsilon}{1+\epsilon}.
\end{aligned}$$

Hence, by (2.5.9),

$$\|H_y(z)^{-1}\|_y \le \frac{1}{(1-(1+\epsilon)\|z-y\|_y)^2}.$$

Since by (2.5.6),

$$\|z-y\|_y \le \frac{\|z-y\|_x}{1-\|y-x\|_x} \le \frac{\|z-y\|_x}{1-(1+\epsilon)\|y-x\|_x},$$

we thus have

$$\begin{aligned}
\|H_y(z)^{-1}\|_y &\le \left(\frac{1-(1+\epsilon)\|y-x\|_x}{1-(1+\epsilon)(\|z-y\|_x+\|y-x\|_x)}\right)^2 \\
&= \left(\frac{1-(1+\epsilon)\|y-x\|_x}{1-(1+\epsilon)\|z-x\|_x}\right)^2,
\end{aligned}$$

the equality relying on the co-linearity of x, y and z. That is, for all $v \neq 0$,

$$\frac{\|v\|_y}{\|v\|_z} \leq \frac{1-(1+\epsilon)\|y-x\|_x}{1-(1+\epsilon)\|z-x\|_x}.$$

Since

$$\frac{\|v\|_x}{\|v\|_y} \leq \frac{1}{1-(1+\epsilon)\|y-x\|_x},$$

it follows that

$$\frac{\|v\|_x}{\|v\|_z} \leq \frac{1}{1-(1+\epsilon)\|z-x\|_x}.$$

In other words,

$$\|H_x(z)^{-1}\|_x \leq \frac{1}{(1-(1+\epsilon)\|z-x\|_z)^2},$$

completing the proof of the proposition. □

The following theorem provides various equivalent definitions of self-concordant functionals.

Theorem 2.5.11 *The following conditions on a functional f are equivalent:*

- *1a: For all $x, y \in D_f$, if $\|y-x\|_x < 1$ then*

$$\frac{\|v\|_y}{\|v\|_x} \leq \frac{1}{1-\|y-x\|_x} \quad \textit{for all } v \neq 0.$$

- *1b: For all $x \in D_f$, and for all y in some open neighborhood of x,*

$$\frac{\|v\|_y}{\|v\|_x} \leq \frac{1}{1-\|y-x\|_x} \quad \textit{for all } v \neq 0.$$

- *1c: For all $x \in D_f$,*

$$\limsup_{y \to x} \frac{\|I - H_x(y)\|_x}{\|y-x\|_x} \leq 2.$$

Moreover, if f satisfies any one (and hence all) of the above conditions, as well as any one condition from the following list, then f satisfies all conditions from the following list.

- *2a: For all $x \in D_f$ we have $B_x(x, 1) \subseteq D_f$.*

- *2b: There exists* $0 < r \leq 1$ *such that for all* $x \in D_f$ *we have* $B_x(x, r) \subseteq D_f$.
- *2c: If a sequence* $\{x_k\}$ *converges to a point in the boundary* ∂D_f *then* $f(x_k) \to \infty$.

Hence, since SC consists precisely of those functionals satisfying conditions 1a and 2a, by choosing one condition from the first set and one from the second, the set SC can be defined as the set of functionals satisfying the two chosen conditions.

Proof: To prove the theorem, we first establish the equivalence of conditions 1a, 1b and 1c. We then prove conditions 1a and 2b together imply 2a as do conditions 1a and 2c. Trivially, 2a implies 2b. To conclude the proof, it then suffices to recall that by Proposition 2.2.12, conditions 1a and 2a together imply 2c.

Now to establish the equivalence of 1a, 1b and 1c. Trivially, 1a implies 1b. Next note

$$\begin{aligned} \|I - H_x(y)\|_x &= \limsup_{v \neq 0} \left| \frac{\langle v, [I - H_x(y)]v \rangle_x}{\|v\|_x^2} \right| \\ &= \limsup_{v \neq 0} \left| 1 - \frac{\|v\|_y^2}{\|v\|_x^2} \right|. \end{aligned}$$

Consequently, condition 1b implies for y near x,

$$\begin{aligned} \|I - H_x(y)\|_x &\leq \frac{1}{(1 - \|y - x\|_x)^2} - 1 \\ &= \frac{2\|y - x\|_x - \|y - x\|_x^2}{(1 - \|y - x\|_x)^2}. \end{aligned}$$

Thus, 1b implies 1c.

To conclude the proof of the equivalence of 1a, 1b and 1c, we assume 1c holds but 1a does not, then obtain a contradiction.

Since

$$\|H_x(y)\|_x = \sup_{v \neq 0} \frac{\|v\|_y^2}{\|v\|_x^2},$$

condition 1a not holding implies there exist x, y, and $\epsilon > 0$ such that $\|y - x\|_x < \frac{1}{1+\epsilon}$ and

$$\|H_x(y)\|_x > \frac{1}{(1 - (1 + \epsilon)\|y - x\|_x)^2}.$$

Considering points on the line segment between x and y and relying on continuity of the Hessian, it then readily follows that there exists y (possibly $y = x$) satisfying $\|y - x\|_x < \frac{1}{1+\epsilon}$,

$$\|H_x(y)\|_x = \frac{1}{(1-(1+\epsilon)\|y-x\|_x)^2}, \tag{2.5.12}$$

and

$$\|H_x(z)\|_x > \frac{1}{(1-(1+\epsilon)\|z-x\|_x)^2} \tag{2.5.13}$$

for all $z := y + t(y - x)$ where $t > 0$ is sufficiently small.

Condition 1c implies for z near y,

$$\|H_y(z)\|_y \leq \frac{1}{(1-(1+\epsilon)\|z-y\|_y)^2}.$$

That is, for all $v \neq 0$,

$$\frac{\|v\|_z}{\|v\|_y} \leq \frac{1}{(1-(1+\epsilon)\|z-y\|_y)^2}. \tag{2.5.14}$$

Likewise, from (2.5.12),

$$\frac{\|v\|_y}{\|v\|_x} \leq \frac{1}{(1-(1+\epsilon)\|y-x\|_x)^2}. \tag{2.5.15}$$

In particular,

$$\|z-y\|_y \leq \frac{\|z-y\|_x}{1-(1+\epsilon)\|y-x\|_x}.$$

Substituting into (2.5.14) and relying on the co-linearity of x, y and z, we have

$$\frac{\|v\|_z}{\|v\|_y} \leq \frac{1-(1+\epsilon)\|y-x\|_x}{1-(1+\epsilon)\|z-x\|_x}. \tag{2.5.16}$$

From (2.5.15) and (2.5.16), for all $v \neq 0$,

$$\frac{\|v\|_z}{\|v\|_x} \leq \frac{1}{1-(1+\epsilon)\|z-x\|_x},$$

that is,

$$\|H_x(z)\|_x \leq \frac{1}{(1-(1+\epsilon)\|z-x\|_x)^2},$$

contradicting (2.5.13). We have thus proven the equivalence of conditions 1a, 1b and 1c.

Now we prove conditions 1a and 2b together imply 2a. Let $0 < r \leq 1$ be as in 2b, i.e., $B_x(x, r) \subseteq D_f$ for all $x \in D_f$. Assuming $y \in D_f$ satisfies $\|y - x\|_x < 1$ and letting

$$t := r(1 - \|y - x\|_x), \quad z := y + t(y - x),$$

we show 1a and 2b together imply $z \in D_f$. Condition 2a readily follows.

By condition 1a,

$$\|y - x\|_y \leq \frac{\|y - x\|_x}{1 - \|y - x\|_x},$$

from which it follows that $z \in B_y(y, r)$. Hence, by 2b we have $z \in D_f$ as desired.

To conclude the proof of the theorem, it only remains to prove that conditions 1a and 2c together imply 2a.

Assuming $y \in D_f$ satisfies $\|y - x\|_x < 1$, condition 1a implies

$$\begin{aligned} f(y) &= f(x) + \langle g_x(x), y - x \rangle_x + \int_0^1 \int_0^t \langle y - x, H_x(x + s(y - x))(y - x) \rangle_x \, ds \, dt \\ &\leq f(x) + \|g_x(x)\|_x + \int_0^1 \int_0^t \|H_x(x + s(y - x))\|_x \, ds \, dt \\ &\leq f(x) + \|g_x(x)\|_x + \frac{1}{2(1 - \|y - x\|_x)^2}. \end{aligned}$$

In particular, $f(y)$ is bounded away from ∞. By condition 2c we conclude $B_x(x, 1) \subseteq D_f$. □

We turn to the original definition of self-concordancy, due to Nesterov and Nemirovskii. First, some motivation.

We know that if one restricts a self-concordant functional f to subspaces – or translates thereof – one obtains self-concordant functionals. In particular, if f is restricted to a line $t \mapsto x + td$ (where $x, d \in \Re^n$) then

$$\phi(t) := f(x + td)$$

is a univariate self-concordant functional. Since for ϕ we have

$$\|v\|_t = \sqrt{\phi''(t)}|v|,$$

the property

$$\frac{\|v\|_s}{\|v\|_t} \leq \frac{1}{1 - \|s - t\|_t}$$

is identical to

$$\frac{\sqrt{\phi''(s)}}{\sqrt{\phi''(t)}} \leq \frac{1}{1-\sqrt{\phi''(t)}|s-t|}.$$

Squaring both sides, then subtracting 1 from both sides, and finally multiplying both sides by $\phi''(t)/|s-t|$, we find

$$\frac{\phi''(s)-\phi''(t)}{|s-t|} \leq \frac{2\phi''(t)^{3/2}-\phi''(t)^2|s-t|}{(1-\sqrt{\phi''(t)}|s-t|)^2}.$$

If f – and hence ϕ – is thrice-differentiable this implies

$$\phi'''(t) \leq 2\phi''(t)^{3/2}. \tag{2.5.17}$$

This result has a converse. The converse, given by the following theorem, coincides with the original definition of self-concordancy due to Nesterov and Nemirovskii.

Theorem 2.5.18 *Assume $f \in C^3$ and assume each of the univariate functionals ϕ obtained by restricting f to lines intersecting D_f satisfy (2.5.17) for all t in their domains. Furthermore, assume that if a sequence $\{x_k\}$ converges to a point in the boundary ∂D_f then $f(x_k) \to \infty$. Then $f \in \mathcal{SC}$.*

Proof: The proof assumes the reader to be familar with certain properties of differentials.

To prove the theorem, it suffices to prove f satisfies condition 1c and 2c of Theorem 2.5.11. Of course 2c is satisfied by assumption.

Assuming the third differential $D^3(x)$ of f at x is written in terms of a basis which is orthonormal w.r.t. $\langle\ ,\ \rangle_x$, The inequality (2.5.17) is equivalent to requiring

$$|D^3(x)[u,u,u]| \leq 2 \text{ whenever } \|u\|_x \leq 1. \tag{2.5.19}$$

On the other hand, condition 1c is equivalent to requiring

$$|D^3(x)[u,v,v]| \leq 2 \text{ whenever } \|u\|_x, \|v\|_x \leq 1. \tag{2.5.20}$$

However, for any C^3-functional f and for any inner product norm $\|\ \| := \langle\ ,\ \rangle^{1/2}$,

$$\max\{|D^3(x)[u,v,w]| : \|u\|, \|v\|, \|w\| \leq 1\} = \max\{D^3(x)[u,u,u]| : \|u\| \leq 1\}.$$

Hence (2.5.20) follows from (2.5.19). □

The original definition of (strongly non-degenerate) self-concordancy was that f satisfy the assumptions of Theorem 2.5.18. The theorem shows such f to be self-concordant according to the definition relied on in these lecture notes. The definition in these notes is ever-so-slightly less-restrictive by requiring only $f \in C^2$, not $f \in C^3$. For example, letting $\| \;\|$ denote the Euclidean norm, the functionals

$$f(x) := \tfrac{1}{2}\|x\|^2 + \tfrac{1}{6}\|x\|^3$$

and

$$f(x) := e^{\|x\|} - \|x\|$$

are self-concordant according to the definition in these notes, but not according to the original definition. (Neither functional is thrice-differentiable at the origin.)

The definition in these notes was not chosen for the slightly broader set of functionals it defines. It was chosen because it provides the reader upfront with some sense of the geometry underlying self-concordancy, and because it is handy in developing the theory. Nonetheless, the original definition has distinct advantages, especially in proving a functional to be self-concordant. For example, assume $D \subseteq \Re^n$ is open and convex, and assume $F \in C^3$ is a functional which takes on only positive values in D, and only the value 0 on the boundary ∂D. Furthermore, assume that for each line intersecting D, the univariate functional $\phi(t) := f(x + td)$ obtained by restricting F to the line happens to be a polynomial – moreover, a polynomial with only real roots. Then, relying on the original definition of self-concordancy, it is not difficult to prove the functional

$$f(x) := -\ln(F(x))$$

to be self-concordant. For, letting $r_1, \ldots, r_d$ denote the roots of ϕ and assuming w.l.o.g. that ϕ is monic, we have

$$\begin{aligned} \phi'''(t) &= -\tfrac{d^3}{dt^3} \ln \prod_i (t - r_i) \\ &= -\tfrac{d^3}{dt^3} \sum_i \ln(t - r_i) \\ &= -2 \sum_i \frac{1}{(t - r_i)^3} \\ &\leq 2 \left(\frac{1}{(t - r_i)^2} \right)^{3/2}, \end{aligned}$$

the inequality due to the relation $\| \ \|_3 \leq \| \ \|_2$ between the 2-norm and the 3-norm on $\Re^d$.

It is an insightful exercise to show that the self-concordancy of the various logarithmic barrier functions are special cases of the result described in the preceding paragraph. Incidentally, functionals F as above are known as "hyperbolic polynomials."

We close this section with a discussion of the qualifying phrase "strongly non-degenerate," which we have suppressed throughout.

Nesterov and Nemirovskii define *self-concordant functionals* (with no qualifiers) as functionals $f \in C^3$, with open and convex domains, satisfying (2.5.17) for the univariate functionals ϕ obtained by restricting f to lines. They define *strongly* self-concordant functionals as having the additional property that $f(x_k) \to \infty$ if the sequence $\{x_k\}$ converges to a point in the boundary ∂D_f. Finally, strongly *non-degenerate* self-concordant functionals are those which satisfy the yet further property that $H(x)$ is pd for all $x \in D_f$.

It is readily proven that self-concordant functionals (thus defined) have psd Hessians.

One might ask if the Nesterov-Nemirovskii definition of, say, strong self-concordancy has a geometric analogue similar to the definition of strongly non-degenerate self-concordancy used in these lecture notes. One would not expect the analogue to be as simple as the definition in these notes because the bilinear forms

$$\langle u, v \rangle_x := \langle u, H(x)v \rangle$$

need not be inner products if f is not non-degenerate. However, there is indeed an analogue, obtained as a direct extension of the definition for strongly non-degenerate self-concordancy. Roughly, strongly self-concordant functionals are those obtained by extending strongly non-degenerate self-concordant functionals to larger vector spaces by having the functional be constant on parallel slices. Specifically, one can prove (as is done in [1]) that f is strongly self-concordant iff $\Re^n$ can be written as a direct sum $L_1 \oplus L_2$ for which there exists a strongly non-degenerate self-concordant functional h, with $D_h \subseteq L_1$, satisfying $f(x_1, x_2) = h(x_1)$. For example, $f(x) := -\ln(x_1)$ is a strongly self-concordant functional with domain the half-space $\Re_+ \oplus \Re^{n-1}$ in $\Re^n$, but it is not non-degenerate.

If self-concordant (resp. strongly self-concordant) functionals are added, the resulting functional is self-concordant (resp. strongly self-concordant). If one of the summands is strongly non-degenerate, so is the sum. This is an indication of how the theory of self-concordant functionals, and strongly self-concordant functionals, parallels the theory developed in these notes. To get to the heart of the theory expeditously, these notes focus on strongly non-degenerate self-concordant functionals. Those are by far the most important functionals.

Henceforth, we return to our practice of refering to functionals as self-concordant when, strictly speaking, we mean strongly non-degenerate self-concordant.

Bibliography

[1] Yu.E. Nesterov and A.S. Nemirovskii, *Interior Point Methods in Convex Optimization: Theory and Applications*, SIAM, Philadelphia, 1993.