

Variational analysis of spectral functions simplified

D. Drusvyatskiy ^{*} C. Kempton [†]

Abstract. Spectral functions of symmetric matrices – those depending on matrices only through their eigenvalues – appear often in optimization. A cornerstone variational analytic tool for studying such functions is a formula relating their subdifferentials to the subdifferentials of their diagonal restrictions. This paper presents a new, short, and revealing derivation of this result. We then round off the paper with an illuminating derivation of the second derivative of C^2 -smooth spectral functions, highlighting the underlying geometry. All of our arguments have direct analogues for spectral functions of Hermitian matrices, and for singular value functions of rectangular matrices.

Key words. Eigenvalues, singular values, nonsmooth analysis, proximal mapping, subdifferential, Hessian, group actions.

AMS Subject Classification. *Primary 49J52, 15A18; Secondary 49J53, 49R05, 58D19.*

1 Introduction

This work revolves around *spectral functions*. These are functions on the space of $n \times n$ symmetric matrices \mathbf{S}^n that depend on matrices only through their eigenvalues, that is, functions that are invariant under the action of the orthogonal group by conjugation. Spectral functions can always be written in a composite form $f \circ \lambda$, where f is a permutation-invariant function on \mathbf{R}^n and λ is a mapping assigning to each matrix X the vector of eigenvalues $(\lambda_1(X), \dots, \lambda_n(X))$ in nonincreasing order.

A pervasive theme in the study of such functions is that various variational properties of the permutation-invariant function f are inherited by the induced spectral function $f \circ \lambda$; see e.g. [1–5, 16–18]. Take convexity for example. Supposing

University of Washington, Department of Mathematics, Seattle, WA 98195; Research of Drusvyatskiy and Kempton was partially supported by the AFOSR YIP award FA9550-15-1-0237.

^{*} E-mail: ddrusv@uw.edu; <http://www.math.washington.edu/~ddrusv/>

[†] E-mail: yumiko88@uw.edu;

that f is closed and convex, the main result of [7] shows that the Fenchel conjugate of $f \circ \lambda$ admits the elegant representation

$$(f \circ \lambda)^* = f^* \circ \lambda. \quad (1.1)$$

An immediate conclusion is that $f \circ \lambda$ agrees with its double conjugate and is therefore convex, that is, convexity of f is inherited by the spectral function $f \circ \lambda$. A convenient characterization of the subdifferential $\partial(f \circ \lambda)(X)$ in terms of $\partial f(\lambda(X))$ then readily follows [7, Theorem 3.1] — an important result for optimization specialists.

In a follow up paper [8], Lewis showed that even for nonconvex functions f , the following exact relationship holds:

$$\partial(f \circ \lambda)(X) = \{U(\text{Diag } v)U^T : v \in \partial f(\lambda(X)), U \in \mathcal{O}_X^n\}, \quad (1.2)$$

where

$$\mathcal{O}_X^n := \{U \in \mathcal{O}^n : X = U(\text{Diag } \lambda(X))U^T\}.$$

Here, the symbol \mathcal{O}^n denotes the group of orthogonal matrices and the symbols $\partial(f \circ \lambda)$ and ∂f may refer to the Fréchet, limiting, or Clarke subdifferentials; see e.g. [14] for the relevant definitions. Thus calculating the subdifferential of the spectral function $f \circ \lambda$ on \mathbf{S}^n reduces to computing the subdifferential of the usually much simpler function f on \mathbf{R}^n . For instance, subdifferential computation of the k 'th largest eigenvalue function $X \mapsto \lambda_k(X)$ amounts to analyzing a piecewise polyhedral function, the k 'th order statistic on \mathbf{R}^n [8, Section 9]. Moreover, the subdifferential formula allows one to gauge the underlying geometry of spectral functions, through their “active manifolds” [1], for example.

In striking contrast to the convex case [7], the proof of the general subdifferential formula (1.2) requires much finer tools, and is less immediate to internalize. This paper presents a short, elementary, and revealing derivation of equation (1.2) that is no more involved than its convex counterpart. Here's the basic idea. Consider the *Moreau envelope*

$$f_\alpha(x) := \inf_y \left\{ f(y) + \frac{1}{2\alpha} |x - y|^2 \right\}.$$

Similar notation will be used for the envelope of $f \circ \lambda$. In direct analogy to equation (1.1), we will observe that the Moreau envelope satisfies the equation

$$(f \circ \lambda)_\alpha = f_\alpha \circ \lambda,$$

and derive a convenient formula for the corresponding proximal mapping. The case when f is an indicator function was treated in [2], and the argument presented here is a straightforward adaptation, depending solely on the Theobald–von Neumann inequality [19, 20]. The key observation now is independent of the eigenvalue setting:

membership of a vector v in the proximal or in the Fréchet subdifferential of any function g at a point x is completely determined by the local behavior of the univariate function $\alpha \mapsto g_\alpha(x + \alpha v)$ near the origin. The proof of the subdifferential formula (1.2) quickly flows from there. It is interesting to note that the argument uses very little information about the properties of the eigenvalue map, with the exception of the Theobald–von Neumann inequality. Consequently, it applies equally well in a more general algebraic setting of certain isometric group actions, encompassing also an analogous subdifferential formula for functions of singular values derived in [11, 12, 15]; a discussion can be found in the appendix of the arXiv version of the paper. A different Lie theoretic approach in the convex case appears in [9].

We complete the paper by reconsidering the second-order theory of spectral functions. In [10, 16, 17], the authors derived a formula for the second derivative of a C^2 -smooth spectral function. In its simplest form it reads

$$\nabla^2 F(\text{Diag } a)[B] = \text{Diag}(\nabla^2 f(a)\text{diag}(B)) + \mathcal{A} \circ B,$$

where $\mathcal{A} \circ B$ is the Hadamard product and

$$\mathcal{A}_{ij} = \begin{cases} \frac{\nabla f(a)_i - \nabla f(a)_j}{a_i - a_j} & \text{if } a_i \neq a_j \\ \nabla^2 f(a)_{ii} - \nabla^2 f(a)_{ij} & \text{if } a_i = a_j \end{cases}.$$

This identity is quite mysterious, and its derivation is largely opaque geometrically. In the current work, we provide a transparent derivation, making clear the role of the invariance properties of the gradient graph. To this end, we borrow some ideas from [17], while giving them a geometric interpretation.

The outline of the manuscript is as follows. Section 2 records some basic notation and an important preliminary result about the Moreau envelope (Lemma 2.1). Section 3 contains background material on orthogonally invariant functions. Section 4 describes the derivation of the subdifferential formula and Section 5 focuses on the second-order theory – the main results of the paper.

2 Notation

This section briefly records some basic notation, following closely the monograph [14]. The symbol \mathbf{E} will always denote an Euclidean space (finite-dimensional real inner product space) with inner product $\langle \cdot, \cdot \rangle$ and induced norm $|\cdot|$. A closed ball of radius $\varepsilon > 0$ around a point x will be denoted by $\mathcal{B}_\varepsilon(x)$. The closure and the convex hull of a set Q in \mathbf{E} will be denoted by $\text{cl } Q$ and $\text{conv } Q$, respectively.

Throughout, we will consider functions f on \mathbf{E} taking values in the extended real line $\overline{\mathbf{R}} := \mathbf{R} \cup \{\pm\infty\}$. For such a function f and a point \bar{x} , with $f(\bar{x})$ finite, the *proximal subdifferential* $\partial_p f(\bar{x})$ consists of all vectors $v \in \mathbf{E}$ such that there exists

constants $r > 0$ and $\varepsilon > 0$ satisfying

$$f(x) \geq f(\bar{x}) + \langle v, x - \bar{x} \rangle - \frac{r}{2}|x - \bar{x}|^2 \quad \text{for all } x \in \mathcal{B}_\varepsilon(\bar{x}).$$

Whenever f is C^2 -smooth near \bar{x} , the proximal subdifferential $\partial_p f(\bar{x})$ consists only of the gradient $\nabla f(\bar{x})$. A function f is said to be *prox-bounded* if it majorizes some quadratic function. In particular, all lower-bounded functions are prox-bounded. For prox-bounded functions, the inequality in the definition of the proximal subdifferential can be taken to hold globally at the cost of increasing r [14, Proposition 8.46]. The *Fréchet subdifferential* of f at \bar{x} , denoted $\hat{\partial}f(\bar{x})$, consists of all vectors $v \in \mathbf{E}$ satisfying

$$f(x) \geq f(\bar{x}) + \langle v, x - \bar{x} \rangle + o(|x - \bar{x}|).$$

Here, as usual, $o(|x - \bar{x}|)$ denotes any term satisfying $\frac{o(|x - \bar{x}|)}{|x - \bar{x}|} \rightarrow 0$. Whenever f is C^1 -smooth near \bar{x} , the set $\hat{\partial}f(\bar{x})$ consists only of the gradient $\nabla f(\bar{x})$. The subdifferentials $\partial_p f(\bar{x})$ and $\hat{\partial}f(\bar{x})$ are always convex, while $\hat{\partial}f(\bar{x})$ is also closed. The *limiting subdifferential* of f at \bar{x} , denoted $\partial f(\bar{x})$, consists of all vectors $v \in \mathbf{E}$ so that there exist sequences x_i and $v_i \in \hat{\partial}f(x_i)$ with $(x_i, f(x_i), v_i) \rightarrow (\bar{x}, f(\bar{x}), v)$. The same object arises if the vectors v_i are restricted instead to lie in $\partial_p f(x_i)$ for each index i ; see for example [14, Corollary 8.47]. The *horizon subdifferential*, denoted $\partial^\infty f(\bar{x})$, consists of all limits of $\lambda_i v_i$ for some sequences $v_i \in \partial f(x_i)$ and $\lambda_i \geq 0$ satisfying $x_i \rightarrow \bar{x}$ and $\lambda_i \searrow 0$. This object records horizontal “normals” to the epigraph of the function. For example, f is locally Lipschitz continuous around \bar{x} if and only if the set $\partial^\infty f(\bar{x})$ contains only the zero vector.

The two key constructions at the heart of the paper are defined as follows. Given a function $f: \mathbf{E} \rightarrow \overline{\mathbf{R}}$ and a parameter $\alpha > 0$, the *Moreau envelope* f_α and the *proximal mapping* $P_\alpha f$ are defined by

$$f_\alpha(x) := \inf_{y \in \mathbf{E}} \left\{ f(y) + \frac{1}{2\alpha}|y - x|^2 \right\},$$

$$P_\alpha f(x) := \operatorname{argmin}_{y \in \mathbf{E}} \left\{ f(y) + \frac{1}{2\alpha}|y - x|^2 \right\}.$$

Extending the definition slightly, we will set $f_0(x) := f(x)$. It is easy to see that f is *prox-bounded* if and only if there exists some point $x \in \mathbf{E}$ and a real $\alpha > 0$ satisfying $f_\alpha(x) > -\infty$.

The proximal and Fréchet subdifferentials are conveniently characterized by a differential property of the function $\alpha \mapsto f_\alpha(x + \alpha v)$. This observation is recorded below. To this end, for any function $\varphi: [0, \infty) \rightarrow \overline{\mathbf{R}}$, the one-sided derivative will be denoted by

$$\varphi'_+(0) := \lim_{\alpha \searrow 0} \frac{\varphi(\alpha) - \varphi(0)}{\alpha}.$$

Lemma 2.1 (Subdifferential and the Moreau envelope).

Consider an lsc, prox-bounded function $f: \mathbf{E} \rightarrow \overline{\mathbf{R}}$, and a point x with $f(x)$ finite. Fix a vector $v \in \mathbf{E}$ and define the function $\varphi: [0, \infty) \rightarrow \overline{\mathbf{R}}$ by setting $\varphi(\alpha) := f_\alpha(x + \alpha v)$. Then the following are true.

(i) The vector v lies in $\hat{\partial}f(x)$ if and only if

$$\varphi'_+(0) = \frac{|v|^2}{2}. \quad (2.1)$$

(ii) The vector v lies in $\partial_p f(x)$ if and only if there exists $\alpha > 0$ satisfying $x \in P_{\alpha} f(x + \alpha v)$, or equivalently

$$\varphi(\alpha) = f(x) + \frac{|v|^2}{2}\alpha.$$

In this case, the equation above continues to hold for all $\tilde{\alpha} \in [0, \alpha]$.

Proof. Claim (ii) is immediate from definitions; see for example [14, Proposition 8.46]. Hence we focus on claim (i). To this end, note first that the inequality

$$\frac{f_\alpha(x + \alpha v) - f(x)}{\alpha} \leq \frac{|v|^2}{2} \quad \text{holds for any } v \in \mathbf{E}. \quad (2.2)$$

Consider now a vector $v \in \hat{\partial}f(x)$ and any sequences $\alpha_i \searrow 0$ and $x_i \in P_{\alpha_i}(x + \alpha_i v)$. We may assume $x_i \neq x$ since otherwise there's nothing to prove. Clearly x_i tend to x and hence

$$\begin{aligned} f_{\alpha_i}(x + \alpha_i v) - f(x) &= f(x_i) - f(x) + \frac{1}{2\alpha_i} |(x_i - x) - \alpha_i v|^2 \\ &\geq o(|x_i - x|) + \frac{1}{2\alpha_i} |x_i - x|^2 + \frac{\alpha_i}{2} |v|^2. \end{aligned}$$

Consequently, we obtain the inequality

$$\frac{f_{\alpha_i}(x + \alpha_i v) - f(x)}{\alpha_i} \geq \frac{|x_i - x|}{\alpha_i} \cdot \frac{o(|x_i - x|)}{|x_i - x|} + \frac{1}{2} \left| \frac{x_i - x}{\alpha_i} \right|^2 + \frac{|v|^2}{2}.$$

Taking into account (2.2) yields the inequality

$$0 \geq \frac{|x_i - x|}{\alpha_i} \cdot \left(\frac{o(|x_i - x|)}{|x_i - x|} + \frac{1}{2} \left| \frac{x_i - x}{\alpha_i} \right| \right).$$

In particular, we deduce $\frac{x_i - x}{\alpha_i} \rightarrow 0$, and the equation (2.1) follows.

Conversely suppose that equation (2.1) holds, and for the sake of contradiction that v does not lie in $\hat{\partial}f(x)$. Then there exists $\kappa > 0$ and a sequence $y_i \rightarrow x$ satisfying

$$f(y_i) - f(x) - \langle v, y_i - x \rangle \leq -\kappa |y_i - x|.$$

Then for any $\alpha > 0$, observe

$$\begin{aligned} \frac{f_\alpha(x + \alpha v) - f(x)}{\alpha} &\leq \frac{1}{\alpha} \left(f(y_i) - f(x) + \frac{1}{2\alpha} |(y_i - x) - \alpha v|^2 \right) \\ &\leq -\kappa \frac{|y_i - x|}{\alpha} + \frac{1}{2} \left| \frac{y_i - x}{\alpha} \right|^2 + \frac{|v|^2}{2}. \end{aligned}$$

Setting $\alpha_i := \frac{|y_i - x|}{\kappa}$ and letting i tend to ∞ yields a contradiction. \square

3 Symmetry and orthogonal invariance

Next we recall a basic correspondence between symmetric functions and spectral functions of symmetric matrices. The discussion follows that of [8]. Henceforth \mathbf{R}^n will denote an n -dimensional real Euclidean space with a specified basis. Hence one can associate \mathbf{R}^n with a collection of n -tuples (x_1, \dots, x_n) , in which case the inner product $\langle \cdot, \cdot \rangle$ is the usual dot product. The finite group of coordinate permutations of \mathbf{R}^n will be denoted by Π^n . A function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ is *symmetric* whenever it is Π^n -invariant, meaning

$$f(\pi x) = f(x) \quad \text{for all } x \in \mathbf{R}^n \text{ and } \pi \in \Pi^n.$$

It is immediate to verify that if f is symmetric, then so is the Moreau envelope f_α for any $\alpha \geq 0$. This elementary observation will be important later.

The vector space of real $n \times n$ symmetric matrices will be denoted by \mathbf{S}^n and will be endowed with the trace inner product $\langle X, Y \rangle = \text{tr } XY$, and the induced Frobenius norm $|X| = \sqrt{\text{tr } X^2}$. For any $x \in \mathbf{R}^n$, the symbol $\text{Diag } x$ will denote the $n \times n$ matrix with x on its diagonal and with zeros off the diagonal, while for a matrix $X \in \mathbf{S}^n$, the symbol $\text{diag } X$ will denote the n -vector of its diagonal entries.

The group of real $n \times n$ orthogonal matrices will be written as \mathcal{O}^n . The eigenvalue mapping $\lambda: \mathbf{S}^n \rightarrow \mathbf{R}^n$ assigns to each matrix X in \mathbf{S}^n the vector of its eigenvalues $(\lambda_1(X), \dots, \lambda_n(X))$ in a nonincreasing order. A function $F: \mathbf{S}^n \rightarrow \overline{\mathbf{R}}$ is *spectral* if it is \mathcal{O}^n -invariant under the conjugation action, meaning

$$F(UXU^T) = F(X) \quad \text{for all } X \in \mathbf{S}^n \text{ and } U \in \mathcal{O}^n.$$

In other words, spectral functions are those that depend on matrices only through their eigenvalues. A basic fact is that any spectral function F on \mathbf{S}^n can be written as a composition of $F = f \circ \lambda$ for some symmetric function f on \mathbf{R}^n . Indeed, f can be realized as the restriction of F to diagonal matrices $f(x) = F(\text{Diag } x)$.

Two matrices X and Y in \mathbf{S}^n are said to admit a *simultaneous spectral decomposition* if there exists an orthogonal matrix $U \in \mathcal{O}^n$ such that UXU^T and UYU^T are both diagonal matrices. It is well-known that this condition holds if and only

if X and Y commute. The matrices X and Y are said to admit a *simultaneous ordered spectral decomposition* if there exists an orthogonal matrix $U \in \mathcal{O}^n$ satisfying $UXU^T = \text{Diag } \lambda(X)$ and $UYU^T = \text{Diag } \lambda(Y)$. The following result characterizing this property, essentially due to Theobald [19] and von Neumann [20], plays a central role in spectral variation analysis.

Theorem 3.1 (Von Neumann-Theobald). *Any two matrices X and Y in \mathbf{S}^n satisfy the inequality*

$$|\lambda(X) - \lambda(Y)| \leq |X - Y|.$$

Equality holds if and only if X and Y admit a simultaneous ordered spectral decomposition.

This result is often called a trace inequality, since the eigenvalue mapping being 1-Lipschitz (as in the statement above) is equivalent to the inequality

$$\langle \lambda(X), \lambda(Y) \rangle \geq \langle X, Y \rangle \quad \text{for all } X, Y \in \mathbf{S}^n.$$

4 Derivation of the subdifferential formula

In this section, we derive the subdifferential formula for spectral functions. In what follows, for any matrix $X \in \mathbf{S}^n$ define the diagonalizing matrix set

$$\mathcal{O}_X := \{U \in \mathcal{O}^n : U(\text{Diag } \lambda(X))U^T = X\}.$$

The spectral subdifferential formula readily follows from Lemma 2.1 and the following intuitive proposition, a proof of which can essentially be seen in [2, Proposition 8].

Theorem 4.1 (Proximal analysis of spectral functions).

Consider a symmetric function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$. Then the equation

$$(f \circ \lambda)_\alpha = f_\alpha \circ \lambda \quad \text{holds.} \quad (4.1)$$

In addition, the proximal mapping admits the representation:

$$P_\alpha(f \circ \lambda)(X) = \{U(\text{Diag } y)U^T : y \in P_\alpha f(\lambda(X)), U \in \mathcal{O}_X\}. \quad (4.2)$$

Moreover, for any $Y \in P_\alpha(f \circ \lambda)(X)$ the matrices X and Y admit a simultaneous ordered spectral decomposition.

Proof. For any X and Y , applying the trace inequality (Theorem 3.1), we deduce

$$f(\lambda(Y)) + \frac{1}{2\alpha}|Y - X|^2 \geq f(\lambda(Y)) + \frac{1}{2\alpha}|\lambda(Y) - \lambda(X)|^2 \geq f_\alpha(\lambda(X)). \quad (4.3)$$

Taking the infimum over Y , we deduce $(f \circ \lambda)_\alpha(X) \geq f_\alpha(\lambda(X))$. On the other hand, for any $U \in \mathcal{O}_X$, the inequalities hold:

$$\begin{aligned} (f \circ \lambda)_\alpha(X) &= \inf_Y \left\{ f(\lambda(Y)) + \frac{1}{2\alpha} |Y - X|^2 \right\} \\ &= \inf_Y \left\{ f(\lambda(Y)) + \frac{1}{2\alpha} |U^T Y U - \text{Diag } \lambda(X)|^2 \right\} \leq f_\alpha(\lambda(X)). \end{aligned}$$

This establishes (4.1).

To establish equation (4.2), consider first a matrix $U \in \mathcal{O}_X$ and a vector $y \in P_\alpha f(\lambda(X))$, and define $Y := U(\text{Diag } y)U^T$. Then we have

$$(f \circ \lambda)(Y) + \frac{1}{2\alpha} |Y - X|^2 = f(y) + \frac{1}{2\alpha} |y - \lambda(X)|^2 = f_\alpha(\lambda(X)) = (f \circ \lambda)_\alpha(X).$$

Hence the inclusion $Y \in P_\alpha(f \circ \lambda)(X)$ is valid, as claimed. Conversely, fix any matrix $Y \in P_\alpha(f \circ \lambda)(X)$. Then plugging in Y into (4.3), the left-hand-side equals $(f \circ \lambda)_\alpha(X)$ and hence the two inequalities in (4.3) hold as equalities. The second equality immediately yields the inclusion $\lambda(Y) \in P_\alpha f(\lambda(X))$, while the first along with Theorem 3.1 implies that X and Y admit a simultaneous ordered spectral decomposition, as claimed. \square

Combining Lemma 2.1 and Theorem 4.1, the main result of the paper readily follows.

Theorem 4.2 (Subdifferentials of spectral functions). *Consider an lsc symmetric function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$. Then the following equation holds:*

$$\partial(f \circ \lambda)(X) = \{U(\text{Diag } v)U^T : v \in \partial f(\lambda(X)), U \in \mathcal{O}_X\}. \quad (4.4)$$

Analogous formulas hold for the proximal, Fréchet, and horizon subdifferentials.

Proof. Fix a matrix X in the domain of $f \circ \lambda$ and define $x := \lambda(X)$. Without loss of generality, suppose that f is lower-bounded. Indeed if this were not the case, then since f is lsc there exists $\varepsilon > 0$ so that f is lower-bounded on the ball $\mathcal{B}_\varepsilon(x)$. Consequently adding to f the indicator function of the symmetric set $\cup_{\pi \in \Pi} \mathcal{B}_\varepsilon(\pi x)$ assures that the function is lower-bounded.

We first dispense with the easy inclusion \subseteq for all the subdifferentials. To this end, recall that if V is a proximal subgradient of $f \circ \lambda$ at X , then there exists $\alpha > 0$ satisfying $X \in P_\alpha(f \circ \lambda)(X + \alpha V)$. Theorem 4.1 then implies that X and V commute. Taking limits, we deduce that all Fréchet, limiting, and horizon subgradients of $f \circ \lambda$ at X also commute with X . Recalling that commuting matrices admit simultaneous spectral decomposition, basic definitions immediately yield the inclusion \subseteq in equation (4.4) for the proximal and for the Fréchet subdifferentials. Taking limits, we deduce the inclusion \subseteq in (4.4) for the limiting and for the horizon subdifferentials, as well.

Next, we argue the reverse inclusion. To this end, define $V := U(\text{Diag } v)U^T$ for an arbitrary matrix $U \in \mathcal{O}_X$ and any vector $v \in \mathbf{R}^n$. Then Theorem 4.1, along with the symmetry of the envelope f_α , yields the equation

$$\frac{(f \circ \lambda)_\alpha(X + \alpha V) - f(\lambda(X))}{\alpha} = \frac{f_\alpha(x + \alpha v) - f(x)}{\alpha}.$$

Consequently if v lies in $\partial_p f(x)$, then Lemma 2.1 shows that for some $\alpha > 0$ the right-hand-side equals $\frac{|v|^2}{2}$, or equivalently $\frac{|V|^2}{2}$. Lemma 2.1 then yields the inclusion $V \in \partial_p(f \circ \lambda)(X)$. Similarly if v lies in $\hat{\partial} f(x)$, then the same argument but with α tending to 0 shows that V lies in $\hat{\partial}(f \circ \lambda)(X)$. Thus the inclusion \supseteq in equation (4.4) holds for the proximal and for the Fréchet subdifferentials. Taking limits, the same inclusion holds for the limiting and for the horizon subdifferentials. This completes the proof. \square

Remark 4.3. It easily follows from Theorem 4.2 that the inclusion \supseteq holds for the Clarke subdifferential. The reverse inclusion, however, requires a separate argument given in [8, Sections 7-8].

In conclusion, we should mention that all the arguments in the section apply equally well for Hermitian matrices (with the standard Hermitian trace product), with the orthogonal matrices replaced by unitary matrices. Entirely analogous arguments also apply for functions of singular values of rectangular matrices (real or complex). For more details, see the appendix in the arXiv version of the paper.

5 Hessians of C^2 -smooth spectral functions

In this section, we revisit the second-order theory of spectral functions. To this end, fix for the entire section an lsc symmetric function $f: \mathbf{R}^n \rightarrow \overline{\mathbf{R}}$ and define the spectral function $F := f \circ \lambda$ on \mathbf{S}^n . It is well known that f is C^2 -smooth around a matrix X if and only if F is C^2 -smooth around $\lambda(X)$; see [10, 16–18]. Moreover, a formula for the Hessian of F is available: for matrices $A = \text{Diag}(a)$ and $B \in \mathbf{S}^n$ we have

$$\nabla^2 F(A)[B] = \text{Diag}(\nabla^2 f(a)\text{diag}(B)) + \mathcal{A} \circ B,$$

where $\mathcal{A} \circ B$ is the Hadamard product and

$$\mathcal{A}_{ij} = \begin{cases} \frac{\nabla f(a)_i - \nabla f(a)_j}{a_i - a_j} & \text{if } a_i \neq a_j \\ \nabla^2 f(a)_{ii} - \nabla^2 f(a)_{ij} & \text{if } a_i = a_j \end{cases}.$$

The assumption that A is a diagonal matrix is made without loss of generality, as will be apparent shortly. In this section, we provide a transparent geometric derivation of the Hessian formula by considering invariance properties of $\text{gph } \nabla F$. Some of our arguments give a geometric interpretation of the techniques in [17].

Remark 5.1 (Hessian and the gradient graph). Throughout the section we will appeal to the following basic property of the Hessian. For any C^2 -smooth function g on an Euclidean space, the vector $z := \nabla^2 g(a)[b]$ is the unique vector satisfying $(z, -b) \in N_{\text{gph } \nabla g}(a, \nabla g(a))$.

Consider now the action of the orthogonal group \mathcal{O}^n on \mathbf{S}^n by conjugation namely $U.X = UXU^T$. Recall that F is *invariant* under this action, meaning $F(U.X) = F(X)$ for all orthogonal matrices U . This action naturally extends to the product space $\mathbf{S}^n \times \mathbf{S}^n$ by setting $U.(X, Y) = (U.X, U.Y)$. As we have seen, the graph $\text{gph } \nabla F$ is then *invariant* with respect to this action:

$$U.\text{gph } \nabla F = \text{gph } \nabla F \quad \text{for all } U \in \mathcal{O}^n.$$

One immediate observation is that $N_{\text{gph } \nabla F}(U.X, U.Y) = U.N_{\text{gph } \nabla F}(X, Y)$. Consequently we deduce

$$(Z, -B) \in N_{\text{gph } \nabla F}(X, Y) \iff (U.Z, -U.B) \in N_{\text{gph } \nabla F}(U.X, U.Y)$$

The formula

$$\nabla^2 F(X)[B] = U^T.\nabla^2 F(U.X)[U.B] \tag{5.1}$$

now follows directly from Remark 5.1, whenever F is C^2 -smooth around X . As a result, when speaking about the operator $\nabla^2 F(X)$, we may assume without loss of generality that X and $\nabla F(X)$ are both diagonal matrices.

Next we briefly recall a few rudimentary properties of the conjugation action; see for example [6, Sections 4, 8, 9]. We say that a $n \times n$ matrix W is *skew-symmetric* if $W^T = -W$. Then it is well-known that \mathcal{O}^n is a smooth manifold and the tangent space to \mathcal{O}^n at the identity matrix consists of skew-symmetric matrices:

$$T_{\mathcal{O}^n}(I) = \{W \in \mathbf{R}^{n \times n} : W \text{ is skew-symmetric}\}.$$

The *commutator* of two matrices $A, B \in \mathbf{R}^{n \times n}$, denoted by $[A, B]$ is the matrix $[A, B] := AB - BA$. An easy computation shows that the commutator of a skew-symmetric matrix with a symmetric matrix is itself symmetric. Moreover, the identity

$$\langle X, [W, Z] \rangle = \langle [X, W], Z \rangle$$

holds for any matrices $X, Z \in \mathbf{S}^n$ and skew-symmetric W . For any matrix $A \in \mathbf{S}^n$, the *orbit* of A , denoted by $\mathcal{O}^n.A$ is the set

$$\mathcal{O}^n.A = \{U.A : U \in \mathcal{O}^n\}.$$

Similarly, the orbit of a pair $(A, B) \in \mathbf{S}^n \times \mathbf{S}^n$ is the set

$$\mathcal{O}^n.(A, B) = \{(U.A, U.B) : U \in \mathcal{O}^n\}.$$

An standard computation¹ now shows that orbits are smooth manifolds with tangent spaces

$$\begin{aligned} T_{\mathcal{O}^n.A}(A) &= \{[W, A] : W \text{ is skew-symmetric}\}, \\ T_{\mathcal{O}^n.(A,B)}(A, B) &= \{([W, A], [W, B]) : W \text{ is skew-symmetric}\}. \end{aligned}$$

Now supposing that F is twice differentiable at a matrix $A \in \mathbf{S}^{n \times n}$, the graph $\text{gph } \nabla F$ certainly contains the orbit $\mathcal{O}^n.(A, \nabla F(A))$. In particular, this implies that the tangent space to $\text{gph } \nabla F$ at $(A, \nabla F(A))$ contains the tangent space to the orbit:

$$\{([W, A], [W, \nabla F(A)]) : W \text{ skew-symmetric}\}.$$

Thus for any $B \in \mathbf{S}^n$, the tuple $(\nabla^2 F(A)[B], -B)$ is orthogonal to the tuple $([W, A], [W, \nabla F(A)])$ for any skew-symmetric matrix W . We record this elementary observation in the following lemma. This also appears as [17, Lemma 3.2].

Lemma 5.2 (Orthogonality to orbits). *Suppose F is C^2 -smooth around $A \in \mathbf{S}^n$. Then for any skew-symmetric matrix W and any $B \in \mathbf{S}^n$, we have*

$$\langle \nabla^2 F(A)[B], [W, A] \rangle = \langle B, [W, \nabla F(A)] \rangle.$$

Proof. This is immediate from the preceding discussion. □

Next recall that the *stabilizer* of a matrix $A \in \mathbf{S}^n$ is the set:

$$\text{Stab}(A) = \{U \in \mathcal{O}^n : U.A = A\}.$$

Similarly we may define the set $\text{Stab}(A, B)$.

Lemma 5.3 (Tangent space to the stabilizer). *For any matrices $A, B \in \mathbf{S}^n$, the tangent spaces to $\text{Stab}(A)$ and to $\text{Stab}(A, B)$ at the identity matrix are the sets*

$$\begin{aligned} &\{W \in \mathbf{R}^{n \times n} : W \text{ skew-symmetric}, [W, A] = 0\}, \\ &\{W \in \mathbf{R}^{n \times n} : W \text{ skew-symmetric}, [W, A] = [W, B] = 0\}, \end{aligned}$$

respectively.

(Proof sketch). Define the orbit map $\theta^{(A)} : \mathcal{O}^n \rightarrow \mathcal{O}^n.A$ by setting $\theta^{(A)}(U) := U.A$. A quick computation shows that $\theta^{(A)}$ is equivariant with respect to left-multiplication action of \mathcal{O}^n on itself and the conjugation action of \mathcal{O}^n on $\mathcal{O}^n.A$. Hence the equivariant rank theorem ([6, Theorem 7.25]) implies that $\theta^{(A)}$ has constant rank. In fact, since $\theta^{(A)}$ is surjective, it is a submersion. It follows that the stabilizer

$$\text{Stab}(A) = (\theta^{(A)})^{-1}(A)$$

is a smooth manifold with tangent space at the identity equal to the kernel of the differential $d\theta^{(A)}|_{U=I}(W) = [W, A]$. The expression for the tangent space to $\text{Stab}(A)$ immediately follows. The analogous expression for $\text{Stab}(A, B)$ follows along similar lines. □

¹Compute the differential of the mapping $\mathcal{O}^n \ni U \mapsto U.A$

With this, we are able to state and prove the main theorem.

Theorem 5.4 (Hessian of C^2 -smooth spectral functions). *Consider a symmetric function $f: \mathbf{R}^n \rightarrow \mathbf{R}$ and the spectral function $F = f \circ \lambda$. Suppose that F is C^2 -smooth around a matrix $A := \text{Diag}(a)$ and for any matrix matrix $B \in \mathbf{S}^n$ define $Z := \nabla^2 F(A)[B]$. Then equality*

$$\text{diag}(Z) = \nabla^2 f(a)[\text{diag}(B)],$$

holds, while for indices $i \neq j$, we have

$$Z_{ij} = \begin{cases} B_{ij} \left(\frac{\nabla f(a)_i - \nabla f(a)_j}{a_i - a_j} \right) & \text{if } a_i \neq a_j \\ B_{ij} (\nabla^2 f(a)_{ii} - \nabla^2 f(a)_{jj}) & \text{if } a_i = a_j. \end{cases}$$

Proof. First observe that clearly f must be C^2 smooth at a . Now, since A is diagonal, so is the gradient $\nabla F(A)$. So without loss of generality, we can assume $\nabla F(A) = \text{Diag}(\nabla f(a))$.

Observe now that $(Z, -B)$ is orthogonal to the tangent space of $\text{gph } \nabla F$ at $(A, \nabla F(A))$. On the other hand, for any vector $a' \in \mathbf{R}^n$, we have equality

$$\left\langle \begin{pmatrix} Z \\ -B \end{pmatrix}, \begin{pmatrix} \text{Diag}(a') - \text{Diag}(a) \\ \text{Diag}(\nabla f(a')) - \text{Diag}(\nabla f(a)) \end{pmatrix} \right\rangle = \left\langle \begin{pmatrix} \text{diag}(Z) \\ -\text{diag}(B) \end{pmatrix}, \begin{pmatrix} a' - a \\ \nabla f(a') - \nabla f(a) \end{pmatrix} \right\rangle.$$

It follows immediately that the tuple $(\text{diag}(Z), -\text{diag}(B))$ is orthogonal to the tangent space of $\text{gph } \nabla f$ at $(a, \nabla f(a))$. Hence we deduce the equality $\text{diag}(Z) = \nabla^2 f(a)[\text{diag}(B)]$ as claimed.

Next fix indices i and j with $a_i \neq a_j$, and define the skew-symmetric matrix $W^{(i,j)} := e_i e_j^T - e_j e_i^T$, where e_k denotes the k 'th standard basis vector. Applying Lemma 5.2 with the skew-symmetric matrix $W = \frac{1}{a_i - a_j} W^{(i,j)}$, we obtain

$$\begin{aligned} -2Z_{ij} &= \left\langle Z, \left[\frac{1}{a_i - a_j} W^{(i,j)}, A \right] \right\rangle = - \left\langle \left[\frac{1}{a_i - a_j} W^{i,j}, B \right], \nabla F(A) \right\rangle \\ &= - \left\langle \text{diag} \left[\frac{1}{a_i - a_j} W^{i,j}, B \right], \nabla f(a) \right\rangle = -2B_{ij} \left(\frac{\nabla f(a)_i - \nabla f(a)_j}{a_i - a_j} \right). \end{aligned}$$

The claimed formula $Z_{ij} = B_{ij} \left(\frac{\nabla f(a)_i - \nabla f(a)_j}{a_i - a_j} \right)$ follows.

Finally, fix indices i and j , with $a_i = a_j$. Observe now the inclusion

$$\text{Stab}(A) \subset \text{Stab}(\nabla F(A)).$$

Indeed for any matrix $U \in \text{Stab}(A)$, we have

$$\nabla F(A) = \nabla F(UAU^T) = U \nabla F(A) U^T.$$

This in particular immediately implies that the tangent space $T_{\text{gph } \nabla F}(A, \nabla F(A))$ is invariant under the action of $\text{Stab}(A)$, that is

$$U.T_{\text{gph } \nabla F}(A, \nabla F(A)) = T_{\text{gph } \nabla F}(A, \nabla F(A))$$

for any $U \in \text{Stab}(A)$. Hence their entire orbit $\text{Stab}(A).(X, Y)$ of any tangent vector $(X, Y) \in T_{\text{gph } \nabla F}(A, \nabla F(A))$ is contained in the tangent space $T_{\text{gph } \nabla F}(A, \nabla F(A))$. We conclude that the tangent space to such an orbit $\text{Stab}(A).(X, Y)$ at (X, Y) is contained in $T_{\text{gph } \nabla F}(A, \nabla F(A))$ as well.

Define now the matrices $E_i := \text{Diag}(e_i)$ and $\hat{Z} := \text{Diag}(\nabla^2 f(a)[e_i])$. Because F is C^2 -smooth, clearly the inclusion $(E_i, \hat{Z}) \in T_{\text{gph } \nabla F}(A, \nabla F(A))$ holds. The above argument, along with Lemma 5.3, immediately implies the inclusion

$$\{([W, E_i], [W, \hat{Z}]) : W \text{ skew-symmetric}, [W, A] = 0\} \subseteq T_{\text{gph } \nabla F}(A, \nabla F(A))$$

and in particular, $([W, E_i], [W, \hat{Z}])$ is orthogonal to $(Z, -B)$ for any skew-symmetric W satisfying $[W, A] = 0$. To finish the proof, simply set $W = W^{(i,j)}$. Then since $a_i = a_j$, we have $[W, A] = 0$ and therefore

$$\begin{aligned} -2Z_{ij} &= \langle Z, [W^{(i,j)}, E_i] \rangle = \langle B, [W^{(i,j)}, \hat{Z}] \rangle = -\langle [W^{(i,j)}, B], \hat{Z} \rangle \\ &= -2B_{ij}(\nabla^2 f(a)_{ii} - \nabla^2 f(a)_{ij}), \end{aligned}$$

as claimed. This completes the proof. \square

Remark 5.5. The appealing geometric techniques presented in this section seem promising for obtaining at least necessary conditions for the generalized Hessian, in the sense of [13], of spectral functions that are not necessarily C^2 -smooth. Indeed the arguments presented deal entirely with the graph $\text{gph } \nabla f$, a setting perfectly adapted to generalized Hessian computations. There are difficulties, however. To illustrate, consider a matrix $Z \in \partial^2 F(A|V)$. Then one can easily establish properties of $\text{Diag } Z$ analogous to those presented in Theorem 5.4, as well as properties of Z_{ij} for indices i and j satisfying $a_i \neq a_j$. The difficulty occurs for indices i and j with $a_i = a_j$. In this case, our argument used explicitly the fact that tangent cones to $\text{gph } \partial f$ are linear subspaces, a property that is decisively false in the general setting.

References

- [1] A. Daniilidis, D. Drusvyatskiy, and A.S. Lewis. Orthogonal invariance and identifiability. *SIAM J. Matrix Anal. Appl.*, 35(2):580–598, 2014.
- [2] A. Daniilidis, A.S. Lewis, J. Malick, and H. Sendov. Prox-regularity of spectral functions and spectral sets. *J. Convex Anal.*, 15(3):547–560, 2008.

- [3] A. Daniilidis, J. Malick, and H.S. Sendov. Locally symmetric submanifolds lift to spectral manifolds. *Preprint U.A.B. 23/2009, 43 p., arXiv:1212.3936 [math.OA]*, 2012.
- [4] C. Davis. All convex invariant functions of hermitian matrices. *Arch. Math.*, 8:276–278, 1957.
- [5] D. Drusvyatskiy and M. Larsson. Approximating functions on stratified sets. *Trans. Amer. Math. Soc.*, 367(1):725–749, 2015.
- [6] J.M. Lee. *Introduction to smooth manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2013.
- [7] A.S. Lewis. Convex analysis on the Hermitian matrices. *SIAM J. Optim.*, 6(1):164–177, 1996.
- [8] A.S. Lewis. Nonsmooth analysis of eigenvalues. *Math. Program.*, 84(1, Ser. A):1–24, 1999.
- [9] A.S. Lewis. Convex analysis on Cartan subspaces. *Nonlinear Anal.*, 42(5, Ser. A: Theory Methods):813–820, 2000.
- [10] A.S. Lewis and H.S. Sendov. Twice differentiable spectral functions. *SIAM J. Matrix Anal. Appl.*, 23(2):368–386 (electronic), 2001.
- [11] A.S. Lewis and H.S. Sendov. Nonsmooth analysis of singular values. I. Theory. *Set-Valued Anal.*, 13(3):213–241, 2005.
- [12] A.S. Lewis and H.S. Sendov. Nonsmooth analysis of singular values. II. Applications. *Set-Valued Anal.*, 13(3):243–264, 2005.
- [13] B.S. Mordukhovich. *Variational Analysis and Generalized Differentiation I: Basic Theory*. Grundlehren der mathematischen Wissenschaften, Vol 330, Springer, Berlin, 2006.
- [14] R.T. Rockafellar and R.J-B. Wets. *Variational Analysis*. Grundlehren der mathematischen Wissenschaften, Vol 317, Springer, Berlin, 1998.
- [15] H.S. Sendov. *Variational spectral analysis*. ProQuest LLC, Ann Arbor, MI, 2001. Thesis (Ph.D.)–University of Waterloo (Canada).
- [16] H.S. Sendov. The higher-order derivatives of spectral functions. *Linear Algebra Appl.*, 424(1):240–281, 2007.
- [17] M. Šilhavý. Differentiability properties of isotropic functions. *Duke Math. J.*, 104(3):367–373, 2000.

- [18] J. Sylvester. On the differentiability of $O(n)$ invariant functions of symmetric matrices. *Duke Math. J.*, 52(2):475–483, 1985.
- [19] C.M. Theobald. An inequality for the trace of the product of two symmetric matrices. *Math. Proc. Cambridge Philos. Soc.*, 77:265–267, 1975.
- [20] J. von Neumann. Some matrix inequalities and metrization of matrix-space. *Tomsk. Univ. Rev.*, 1:286–300, 1937.