MAT 237: Supplementary notes on Lagrange multipliers

In these notes we provide a little more information about Lagrange multipliers than can be found in the basic online notes for MAT237.

Suppose K is a compact subset of $\mathbb{R}^n$. Then if $f\colon K \to \mathbb{R}$ is a continuous function, there exists a point $\mathbf{a}$ in K at which f is minimal (as well as a point where f is maximal). In our discussion of the Hessian matrix, we have seen how to find and characterize such a point if $\mathbf{a}$ lies in the *interior* of K. But how do we find $\mathbf{a}$ if it lies on the *boundary* of K? Such a problem is more subtle, and here we will consider the case where $S = \partial K$ is defined as the level set of a continuously differentiable function. More specifically, for $m < n$ let $\mathbf{G}\colon \mathbb{R}^n \to \mathbb{R}^m$ be C^1, and suppose that at every point $\mathbf{x} \in \mathbf{G}^{-1}(\mathbf{0})$ the matrix $D\mathbf{G}(\mathbf{x})$ has rank m (in which case we say $\mathbf{0}$ is a *regular value* of $\mathbf{G}$). We will consider the problem of minimizing (or maximizing) a C^1 function $f\colon \mathbb{R}^n \to \mathbb{R}$ on the set $S = \mathbf{G}^{-1}(\mathbf{0})$. In applications, this is known as a *constrained optimization* problem, where the goal is optimizing the function f subject to the constraints $G_k(\mathbf{x}) = 0$ for $k = 1, \ldots, m$, where $\mathbf{G}(\mathbf{x}) = (G_1(\mathbf{x}), \ldots, G_m(\mathbf{x}))$. We fix the notation used in this paragraph throughout these notes.

Remark 1. More generally, we could consider the case where S is an embedded submanifold of $\mathbb{R}^n$. The situation above is a special case of this thanks to the *Implicit Function Theorem*, which guarantees that the pre-image of a regular value of $\mathbf{G}$ is an embedded $(n-m)$-dimensional submanifold of $\mathbb{R}^n$. We will discuss these notions later in the course, but for now we avoid this terminology.

Before proving the main theorems associated with constrained optimization, let us study a simple example to develop an idea of how best to proceed.

Example 1. Let us set $f(x, y) = xy$, $m = 1$, and $\mathbf{G}(x, y) = G(x, y) = x^2 + 2y^2 - 5$. That is, we will find the minimal values of the function $f(x, y) = xy$ on the ellipse $x^2 + 2y^2 = 5$. Note that $\nabla G(x, y)$ is never zero on the ellipse, so that 0 is indeed a regular value of G. To see how to solve this problem, let us analyse what happens to the level curves $f(x, y) = c$ as c varies.

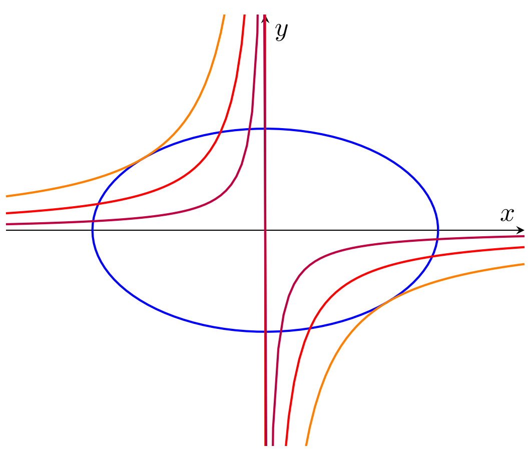

FIGURE 1. Level sets of $f(x, y) = xy$.

The level sets $c = -1/4$, $c = -1/\sqrt{2}$, and $c = -\sqrt{2}$ are plotted in Figure 1 above. We see that as $|c|$ increases, the level sets move further and further away from the origin. It stands to reason that the smallest value of c we can obtain while still having the level set $f(x, y) = c$ intersect the ellipse will occur when the level set is exactly *tangent* to the ellipse. Since ∇f is perpendicular to the level sets of f, such a point of tangency will occur when $\nabla f(x, y)$ is perpendicular to the ellipse. But ∇G is perpendicular to the ellipse at every point, so we see that $\nabla f(x, y)$ and $\nabla G(x, y)$ will be *parallel* at the point of tangency. In other words, there exists some $\lambda \in \mathbb{R}$ such that $\nabla f(x, y) = \lambda \nabla G(x, y)$. Thus in order to solve the problem, we should find all points (x, y) where this occurs.

First we must compute ∇f and ∇G:

$$\nabla f(x, y) = y\mathbf{e}_1 + x\mathbf{e}_2, \qquad \nabla G(x, y) = 2x\mathbf{e}_1 + 4y\mathbf{e}_2.$$

Introducing a parameter λ and setting $\nabla f(x, y) = \lambda \nabla G(x, y)$ yields the system of equations

$$\begin{aligned} y &= 2\lambda x \\ x &= 4\lambda y. \end{aligned}$$

Substituting the first equation into the second gives $1 = 8\lambda^2$, or $\lambda = \pm\frac{1}{2\sqrt{2}}$ which gives $y = \pm\frac{x}{\sqrt{2}}$. Plugging this in to the equation $x^2+2y^2 = 4$, we find that $x = \pm\sqrt{2}$ and hence $y = \pm 1$. We conclude that the minimum value of f on the ellipse is given by $-\sqrt{2}$, which occurs at the points $(-\sqrt{2}, 1)$ and $(\sqrt{2}, -1)$.

The key to solving the example above was to look for points (x, y) on $G^{-1}(0)$ where $\nabla f(x, y)$ was orthogonal to the tangent line to $G^{-1}(0)$ at (x, y). In order to generalize this approach to functions on $\mathbb{R}^n$, we need a higher-dimensional analogue of the notion of tangent lines to curves. When we consider curves in $\mathbb{R}^2$ as in the example above, there is only one possibility for the tangent line to the curve at a point. On the other hand, for a two-dimensional surface in $\mathbb{R}^3$, there are many lines tangent to the surface at a given point. In Figure 2 below, we see that there is an entire *plane* in $\mathbb{R}^3$ tangent to the sphere S^2 at the north pole.

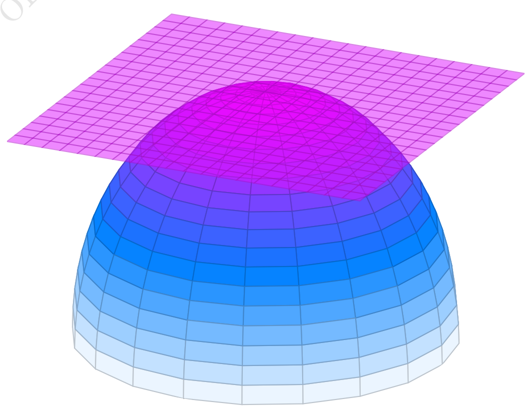

FIGURE 2. Surfaces have tangent *planes*.

In general, for an $(n-1)$-dimensional hypersurface S in $\mathbb{R}^n$ and a point $\mathbf{x} \in S$, the set of vectors in $\mathbb{R}^n$ tangent to S at $\mathbf{x}$ is an $(n-1)$-dimensional vector space.

We denote this vector space $T_{\mathbf{x}}S$ and call it the *tangent space* to S at $\mathbf{x}$. The most geometric way to construct $T_{\mathbf{x}}S$ is to use curves on surfaces. We think of a curve on a surface S as a suitably differentiable function $\gamma\colon J \to S$, where J is some interval containing 0. We say such a curve is *through the point* $\mathbf{x} \in S$ when $\gamma(0) = \mathbf{x}$. When γ is not constant, the image of γ is one-dimensional, so there is only one line tangent to γ at the point $\mathbf{x}$, and $\gamma'(0)$ is parallel to this line. Since the image of γ lies in S, the vector $\gamma'(0)$ is in fact tangent to S. Figure 3 shows a curve (a small circle) on a sphere going through a given point, as well as the tangent vector at that point.

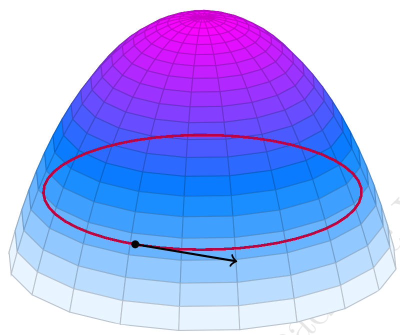

FIGURE 3. Curves on surfaces and tangent vectors.

It seems natural, then, to define $T_{\mathbf{x}}S$ to be the set of all vectors $\mathbf{v} \in \mathbb{R}^n$ such that $\mathbf{v} = \gamma'(0)$ for some curve γ on S through $\mathbf{x}$. We will often think of tangent vectors this way in our computations, and many authors use this definition. The problem is that it is not intrinsic to the surface S. That is, we would like to think of a 2-dimensional sphere S^2, for example, as its own geometric object. But S^2 can be embedded in $\mathbb{R}^n$ for any $n > 2$, and our definition above depends on which $\mathbb{R}^n$ we embed it in. Put another way, we would like to think of the tangent space merely as a real vector space, rather than as a subspace of $\mathbb{R}^n$ for some n. Thus to define $T_{\mathbf{x}}S$, we will use objects that are intrinsic to the surface S, namely, the curves $\gamma\colon J \to S$. We will think of the curves themselves as tangent vectors, with the caveat that two curves γ_1, γ_2 should not be considered distinct if $\gamma_1'(0) = \gamma_2'(0)$. This definition will generalize more readily to the case of manifolds than the definition using vectors in $\mathbb{R}^n$.

Definition 1 (Tangent space)**.** A *curve* in S through $\mathbf{x} \in S$ is a continuous function $\gamma\colon J \to S$, where $J \subseteq \mathbb{R}$ is an open interval containing 0, such that $\gamma(0) = \mathbf{x}$ and γ is differentiable at $t = 0$. Given two such curves, γ_1 and γ_2, we say γ_1 and γ_2 are *equivalent* if $\gamma_1'(0) = \gamma_2'(0)$. This defines an equivalence relation on the set of all curves in S through $\mathbf{x}$. The set of equivalence classes is called the *tangent space* to S at $\mathbf{x}$, and is denoted $T_{\mathbf{x}}S$.

Theorem 1. *Let* $\mathbf{G}\colon \mathbb{R}^n \to \mathbb{R}^m$ *be of class* C^1*, and assume* $\mathbf{0}$ *is a regular value of* $\mathbf{G}$*. Set* $S = \mathbf{G}^{-1}(\mathbf{0})$*. Then for any* $\mathbf{x} \in S$*, the tangent space* $T_{\mathbf{x}}S$ *is an* $(n-m)$*-dimensional real vector space.*

Proof. See, for example, [1, Proposition 4.4.2]. □

Assume we are given a surface S embedded in $\mathbb{R}^n$ for some n, and define $\mathcal{T}_{\mathbf{x}}S = \{\mathbf{v} \mid \mathbf{v} = \gamma'(0) \text{ for some curve } \gamma \text{ in } S \text{ through } \mathbf{x}\}$. Then the map $[\gamma] \mapsto \gamma'(0)$ provides an isomorphism $T_{\mathbf{x}}S \xrightarrow{\sim} \mathcal{T}_{\mathbf{x}}S$, and so there is no harm in thinking of tangent vectors as being elements of $\mathcal{T}_{\mathbf{x}}S$. In what follows, we will make this identification whenever it simplifies matters.

In order to prove the existence of what are called Lagrange multipliers, we need to recall some facts from linear algebra. Given a finite dimensional inner product space $(V, \langle \cdot, \cdot \rangle)$ and a subspace $U \subseteq V$, the *orthogonal complement* of U is defined to be the set

$$U^{\perp} = \{\mathbf{v} \in V \mid \langle \mathbf{v}, \mathbf{u} \rangle = 0 \text{ for all } \mathbf{u} \in U\}.$$

The orthogonal complement of U is always a subspace of V, and

$$\dim U^{\perp} = \dim V - \dim U.$$

It is also worth noting that we always have $(U^{\perp})^{\perp} = U$.

Theorem 2. *Let $f \colon \mathbb{R}^n \to \mathbb{R}$ and $\mathbf{G} \colon \mathbb{R}^n \to \mathbb{R}^m$ be of class C^1. Assume $\mathbf{0}$ is a regular value of $\mathbf{G}$, and set $S = \mathbf{G}^{-1}(\mathbf{0})$. Suppose $\mathbf{a} \in S$ is a local minimum or a local maximum for f in S. Then there exist unique scalars $\lambda_1, \ldots, \lambda_m \in \mathbb{R}$ with $\nabla f(\mathbf{a}) = \sum_{k=1}^{m} \lambda_k \nabla G_k(\mathbf{a})$.*

Proof. Assume $\mathbf{a}$ is a minimum for f in S, the proof for a maximum is similar. Let γ be a curve in S through $\mathbf{a}$, and define the function $h(t) = f(\gamma(t))$. We claim that h has a local minimum at $t = 0$. Indeed, since $\mathbf{a}$ is a local minimum for f in S, there exists $\epsilon > 0$ such that $f(\mathbf{x}) \geq f(\mathbf{a})$ whenever $\mathbf{x} \in S$ and $||\mathbf{x} - \mathbf{a}|| < \epsilon$. Since γ is continuous, there exists $\delta > 0$ such that $||\gamma(t) - \gamma(0)|| < \epsilon$ whenever $|t| < \delta$. But $\gamma(0) = \mathbf{a}$ and $\gamma(t) \in S$ for all t, so we have that $f(\gamma(t)) \geq f(\gamma(0))$ whenever $|t| < \delta$, which shows that h has a local minimum at $t = 0$. In particular, $h'(0) = 0$. On the other hand, by the chain rule we have

$$h'(0) = \nabla f(\gamma(0)) \cdot \gamma'(0) = \nabla f(\mathbf{a}) \cdot \gamma'(0). \tag{1}$$

Since γ was arbitrary, we see that $\nabla f(\mathbf{a}) \cdot \mathbf{v} = 0$ for all $\mathbf{v} \in T_{\mathbf{x}}S$. In other words, $\nabla f(\mathbf{a}) \in (T_{\mathbf{x}}S)^{\perp}$. We also have $\nabla G_k(\mathbf{a}) \in (T_{\mathbf{x}}S)^{\perp}$ for all k (since S is contained in a level set of each G_k). Now $\mathbf{0}$ is a regular value of $\mathbf{G}$, so the vectors $\nabla G_k(\mathbf{a})$ are linearly independent (being the rows of the matrix $D\mathbf{G}(\mathbf{a})$, which has rank m), and since $\dim(T_{\mathbf{x}}S)^{\perp} = m$ they form a basis of $(T_{\mathbf{x}}S)^{\perp}$. By writing the vector $\nabla f(\mathbf{a})$ in this basis, we obtain unique scalars $\lambda_1, \ldots, \lambda_m$ such that

$$\nabla f(\mathbf{a}) = \sum_{k=1}^{m} \lambda_k \nabla G_k(\mathbf{a}). \qquad \square$$

The numbers $\lambda_1, \ldots, \lambda_m$ guaranteed by the theorem above are called *Lagrange multipliers*. We hasten to point out that while Theorem 2 gives necessary conditions for $\mathbf{a}$ to be a local extremum of f on S, the conditions are not sufficient, as the following example shows.

Example 2. We will optimize the function $f(x, y, z) = x^2 - y^2$ subject to the constraint $G(x, y, z) = x^2 + 2y^2 + 3z^2 - 1 = 0$. Note that 0 is a regular value of G, so we may apply Theorem 2. If we solve the equation $\nabla f(x, y, z) = \lambda \nabla G(x, y, z)$ we arrive at the following system of equations:

$$\begin{aligned}2x &= 2\lambda x\\ -2y &= 4\lambda y\\ 0 &= 6\lambda z\\ 1 &= z^2 + 2y^2 + 3z^2.\end{aligned}$$

If $\lambda = 0$ then we must have $x = y = 0$ so that $z = \pm\frac{1}{\sqrt{3}}$. If $\lambda \neq 0$ and $x = 0$, then we see that $\lambda = -\frac{1}{2}$ and $z = 0$, so we must have $y = \pm\frac{1}{\sqrt{2}}$. Lastly if $\lambda \neq 0$ and $x \neq 0$ then we have $\lambda = 1$, $y = z = 0$, and $x = 1 \pm 1$. Thus the possibilities for the extrema of f subject to the constraint $G = 0$ are $(\pm 1, 0, 0)$, $(0, \pm\frac{1}{\sqrt{2}}, 0)$, and $(0, 0, \frac{1}{\sqrt{3}})$. In order to determine the nature of these critical points, we need a version of the second derivative test for constrained optimization. In Example 4 below, we will prove that the points $(0, 0, \pm\frac{1}{\sqrt{3}})$ are neither local maxima nor local minima for f on $G^{-1}(0)$ even though they satisfy the conclusion of Theorem 2.

In the following example we explore the meaning of the Lagrange multipliers $\lambda_1, \ldots, \lambda_m$ in terms of the problem of optimizing f subject to the constraints given by $\mathbf{G}$.

Example 3. Let $\mathbf{a}$ be a local minimum for f on S, and let $\lambda_1, \ldots, \lambda_m$ be the Lagrange multipliers guaranteed to exist by Theorem 2. Suppose we apply a small perturbation to our constraints, that is, we replace the function $\mathbf{G}(\mathbf{x})$ by the function $\mathbf{G}_\mathbf{h}(\mathbf{x}) = \mathbf{G}(\mathbf{x}) - \mathbf{h}$ for some vector $\mathbf{h} = (h_1, \ldots, h_m)$. Let us assume that for all $\mathbf{h}$ in some neighbourhood N of $\mathbf{0}$, $D\mathbf{G}_\mathbf{h}(\mathbf{x})$ has rank m for all $\mathbf{x} \in \mathbf{G}_\mathbf{h}^{-1}(\mathbf{0})$ and the problem of minimizing f subject to the constraints $\mathbf{G}_\mathbf{h}(\mathbf{x}) = 0$ has a solution in some neighbourhood of $\mathbf{a}$. Denote this solution by $\mathbf{a}(\mathbf{h})$ and the corresponding Lagrange multipliers by $\lambda_1(\mathbf{h}), \ldots, \lambda_m(\mathbf{h})$. Assume further that the function $\mathbf{a}(\mathbf{h})\colon \mathbb{R}^m \to \mathbb{R}^n$ is differentiable in N. Define the function $f^*\colon \mathbb{R}^m \to \mathbb{R}$ by $f^*(\mathbf{h}) = f(\mathbf{a}(\mathbf{h}))$. Then f^* represents the optimal value of f subject to the constraints $\mathbf{G}_\mathbf{h}(\mathbf{x}) = \mathbf{0}$ as a function of $\mathbf{h}$.

Differentiating the function f^*, we find

$$\begin{aligned}\nabla f^*(\mathbf{h}) &= \sum_{j=1}^{m} \left(\nabla f(\mathbf{a}(\mathbf{h})) \cdot \frac{\partial}{\partial h_j} \mathbf{a}(\mathbf{h}) \right) \mathbf{e}_j\\ &= \sum_{j=1}^{m} \left(\sum_{k=1}^{m} \lambda_k(\mathbf{h}) \nabla G_k(\mathbf{a}(\mathbf{h})) \cdot \frac{\partial}{\partial h_j} \mathbf{a}(\mathbf{h}) \right) \mathbf{e}_j.\end{aligned}$$

Notice that in the second line above, we have used the fact that $\mathbf{a}(\mathbf{h})$ is optimal, as well as the fact that $D\mathbf{G}_\mathbf{h} = D\mathbf{G}$. Since $\mathbf{a}(\mathbf{h})) \in \mathbf{G}_\mathbf{h}^{-1}(\mathbf{0})$, we have $G_k(\mathbf{a}(\mathbf{h})) = h_k$. Differentiating this expression with respect to h_j, we find

$$\nabla G_k(\mathbf{a}(\mathbf{h})) \cdot \frac{\partial}{\partial h_k} \mathbf{a}(\mathbf{h}) = 1, \quad \nabla G_k(\mathbf{a}(\mathbf{h})) \cdot \frac{\partial}{\partial h_j} \mathbf{a}(\mathbf{h}) = 0, \text{ for } j \neq k.$$

Using these equations, we conclude that

$$\nabla f^*(\mathbf{h}) = \sum_{j=1}^{m} \lambda_j(\mathbf{h}) \mathbf{e}_j.$$

We have arrived at the meaning of the Lagrange multiplier: $\lambda_j = \frac{\partial f^*}{\partial h_j}$. In other words, λ_j tells us how the optimal value of f changes when we change the constraint $G_j(\mathbf{x}) = 0$ slightly.

In light of Theorem 2, any point $\mathbf{a} \in S$ for which there exist scalars $\lambda_1, \dots, \lambda_m$ such that $\nabla f(\mathbf{a}) = \sum_{k=1}^m \lambda_k G_k(\mathbf{a})$ shall be called a *critical point of f on S*. If we define the *Lagrangian function* of such a point $\mathcal{L} \colon \mathbb{R}^n \to \mathbb{R}$ by

$$\mathcal{L}(\mathbf{x}) = f(\mathbf{x}) - \sum_{k=1}^{m} \lambda_k G_k(\mathbf{x}),$$

then we see that Theorem 2 above states that

$$\nabla \mathcal{L}(\mathbf{a}) = 0.$$

In this form, the theorem resembles the first derivative test for unconstrained optimization of the Lagrangian function. This suggests that in order to classify such points, we should study the Hessian matrix of $\mathcal{L}$. Recall that in the unconstrained case, a critical point of f is a local minimum if the Hessian matrix of f is positive definite at that point. It turns out that in the case of constrained optimization, it is enough to ask that the Hessian of the Lagrangian function be positive definite on the tangent space $T_{\mathbf{a}}S$, which we now prove.

Theorem 3. *Let $f \colon \mathbb{R}^n \to \mathbb{R}$ and $\mathbf{G} \colon \mathbb{R}^n \to \mathbb{R}^m$ be C^2 functions. Assume $\mathbf{0}$ is a regular value for $\mathbf{G}$, and set $S = \mathbf{G}^{-1}(\mathbf{0})$. Suppose $\mathbf{a} \in \mathbb{R}^n$ is a critical point of f on S with Lagrangian function $\mathcal{L}$, and denote by $H\mathcal{L}(\mathbf{a})$ the Hessian matrix of $\mathcal{L}$ at $\mathbf{a}$.*

(1) *If $\mathbf{a}$ is a local minimum for f on S, then $\mathbf{v}^T H\mathcal{L}(\mathbf{a})\mathbf{v} \geq 0$ for all $\mathbf{v} \in T_{\mathbf{a}}S$.*
(2) *If $\mathbf{a}$ is a local maximum for f on S, then $\mathbf{v}^T H\mathcal{L}(\mathbf{a})\mathbf{v} \leq 0$ for all $\mathbf{v} \in T_{\mathbf{a}}S$.*
(3) *If $\mathbf{v}^T H\mathcal{L}(\mathbf{a})\mathbf{v} > 0$ for all nonzero $\mathbf{v} \in T_{\mathbf{a}}S$, then $\mathbf{a}$ is a local minimum for f on S.*
(4) *If $\mathbf{v}^T H\mathcal{L}(\mathbf{a})\mathbf{v} < 0$ for all nonzero $\mathbf{v} \in T_{\mathbf{a}}S$, then $\mathbf{a}$ is a local maximum for f on S.*
(5) *If there exist $\mathbf{v}_1, \mathbf{v}_2 \in T_{\mathbf{a}}S$ such that $\mathbf{v}_1^T H\mathcal{L}(\mathbf{a})\mathbf{v}_1 > 0$ and $\mathbf{v}_2^T H\mathcal{L}(\mathbf{a})\mathbf{v}_2 < 0$ then $\mathbf{a}$ is neither a local minimum nor a local maximum for f on S.*

Proof. We will prove the statements about minima, the corresponding results about maxima are proved similarly.

(1) Let $\gamma \colon J \to S$ be a curve in S through $\mathbf{a}$. If $\gamma'(0) = \mathbf{0}$, then we have $\gamma'(0)^T H\mathcal{L}(\mathbf{a})\gamma'(0) = 0$ and we are done, so assume that $\gamma'(0) \neq \mathbf{0}$. Choose a strictly decreasing sequence of real numbers $t_j,\ j \in \mathbb{N}$ such that $t_j \to 0$, and set $\mathbf{x}_j = \gamma(t_j)$. Then by continuity of γ, we have $\mathbf{x}_j \to \mathbf{a}$ as $j \to \infty$. Moreover, since γ is differentiable at $t = 0$, we have

$$\lim_{j\to\infty} \frac{\mathbf{x}_j - \mathbf{a}}{||\mathbf{x}_j - \mathbf{a}||} = \frac{\gamma'(0)}{||\gamma'(0)||}. \tag{2}$$

Since f, G_k are all of class C^2, $\mathcal{L}$ is of class C^2 and we have, by the limit definition of the Hessian,

$$\lim_{j\to\infty} \frac{(\mathbf{x}_j - \mathbf{a})^T H\mathcal{L}(\mathbf{a})(\mathbf{x}_j - \mathbf{a})}{||\mathbf{x}_j - \mathbf{a}||^2} = \lim_{j\to\infty} \frac{\mathcal{L}(\mathbf{x}_j) - \mathcal{L}(\mathbf{a}) - (\mathbf{x}_j - \mathbf{a}) \cdot \nabla\mathcal{L}(\mathbf{a})}{||\mathbf{x}_j - \mathbf{a}||^2}.$$

Using equation (2) and the fact that $\nabla\mathcal{L}(\mathbf{a}) = 0$, we find

$$\begin{aligned}\frac{1}{2}\gamma'(0)^T H\mathcal{L}(\mathbf{a})\gamma'(0) &= ||\gamma'(0)||^2 \lim_{j\to\infty} \frac{\mathcal{L}(\mathbf{x}_j) - \mathcal{L}(\mathbf{a}) - (\mathbf{x}_j - \mathbf{a})\cdot\nabla\mathcal{L}(\mathbf{a})}{||\mathbf{x}_j - \mathbf{a}||^2}.\\ &= ||\gamma'(0)||^2 \lim_{j\to\infty} \frac{\mathcal{L}(\mathbf{x}_j) - \mathcal{L}(\mathbf{a})}{||\mathbf{x}_j - \mathbf{a}||^2}.\end{aligned}$$

Now $\mathbf{a}$ is a local minimum for f on S, and $\mathcal{L}(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in S$. It follows that $\mathbf{a}$ is a local minimum for $\mathcal{L}$ on S, and so for large enough j, we have $\mathcal{L}(\mathbf{x}_j) \geq \mathcal{L}(\mathbf{a})$. This shows that $\gamma'(0)^T H\mathcal{L}(\mathbf{a})\gamma'(0) \geq 0$. Since γ was arbitrary, we see that $\mathbf{v}^T H\mathcal{L}(\mathbf{a})\mathbf{v} \geq 0$ for all $\mathbf{v} \in T_{\mathbf{a}}S$.

(3) Since $\mathcal{L}(\mathbf{x}) = f(\mathbf{x})$ for $\mathbf{x} \in S$, it is enough to prove that $\mathbf{a}$ is a local minimum for $\mathcal{L}$ on S. In fact, we claim that there exist positive numbers m, ϵ such that

$$\mathbf{x} \in S \cap B_\epsilon(\mathbf{a}) \implies \mathcal{L}(\mathbf{x}) \geq \mathcal{L}(\mathbf{a}) + m||\mathbf{x} - \mathbf{a}||^2.$$

Indeed, suppose this is not the case. Then there exist vectors $\mathbf{x}_j$ such that $\mathbf{x}_j \in S \cap B_{1/j}(\mathbf{a})$, $\mathbf{x}_j \neq \mathbf{a}$, and $\mathcal{L}(\mathbf{x}_j) - \mathcal{L}(\mathbf{a}) < \frac{1}{j}||\mathbf{x}_j - \mathbf{a}||^2$ for all j. Notice that the sequence $\mathbf{x}_j$ converges to $\mathbf{a}$. Define unit vectors $\mathbf{y}_j$ by $\mathbf{x}_j = \mathbf{a} + t_j\mathbf{y}_j$ for real numbers $t_j > 0$. The sequence t_j converges to 0, and since the unit sphere is compact, the sequence $\mathbf{y}_j$ has a convergent subsequence by the Bolzano-Weierstrass Theorem. Thus by passing to subsequences if necessary, we may assume the sequence $\mathbf{y}_j$ converges to some $\mathbf{y} \in S^{n-1}$. For any j and any k, by the Mean Value Theorem we can find a number $0 < b_{jk} < 1$ such that

$$G_k(\mathbf{x}_j) - G_k(\mathbf{a}) = t_j\mathbf{y}_j \cdot \nabla G_k(\mathbf{a} + b_{jk}t_j\mathbf{y}_j).$$

On the other hand, since $\mathbf{a}, \mathbf{x}_j \in S$, we have $G_k(\mathbf{x}_j) - G_k(\mathbf{a}) = 0$, so by dividing the equation above by t_j and taking the limit $j \to \infty$ we obtain

$$\mathbf{y} \cdot \nabla G_k(\mathbf{a}) = 0.$$

Since the vectors $\nabla G_k(\mathbf{a})$ span $(T_{\mathbf{a}}S)^\perp$ (see the proof of Theorem 2), it follows that $\mathbf{y} \in T_{\mathbf{a}}S$. By the definition of the Hessian matrix of $\mathcal{L}$ we have

$$\begin{aligned}\frac{1}{2}\mathbf{y}^T H\mathcal{L}(\mathbf{a})\mathbf{y} &= \lim_{j\to\infty} \frac{\mathcal{L}(\mathbf{x}_j) - \mathcal{L}(\mathbf{a}) - \nabla\mathcal{L}(\mathbf{a})\cdot(\mathbf{x}_j - \mathbf{a})}{||\mathbf{x}_j - \mathbf{a}||^2}\\ &= \lim_{j\to\infty} \frac{\mathcal{L}(\mathbf{x}_j) - \mathcal{L}(\mathbf{a})}{||\mathbf{x}_j - \mathbf{a}||^2}\\ &\leq \lim_{j\to\infty} \frac{1}{j}\\ &= 0.\end{aligned}$$

Hence if we have $\mathbf{v}^T H\mathcal{L}(\mathbf{a})\mathbf{v} > 0$ for all nonzero $\mathbf{v} \in T_{\mathbf{a}}S$, then our claim above holds and so $\mathbf{a}$ is a local minimum for f on S.

(5) This is a direct consequence of parts (1) and (2). □

We are now ready to complete example 2.

Example 4. When optimizing the function $f(x, y, z) = x^2 - y^2$ subject to the constraint $S = \{(x, y, z) \mid x^2 + 2y^2 + 3z^2 - 1 = 0\}$, we found the possible extrema to be $(\pm 1, 0, 0)$ with Lagrange multiplier $\lambda = 1$, $(0, \pm\frac{1}{\sqrt{2}}, 0)$ with $\lambda = -\frac{1}{2}$, and $(0, 0, \pm\frac{1}{\sqrt{3}})$ with $\lambda = 0$. We now determine the nature of these critical points. For the points $(\pm 1, 0, 0)$, we have $\mathcal{L}(x, y, z) = f(x, y, z) - G(x, y, z) = -3y^2 - 3z^2 - 1$ and so the Hessian is given by

$$H\mathcal{L}(\pm 1, 0, 0) = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -6 & 0 \\ 0 & 0 & -6 \end{pmatrix}.$$

Now at the points $(\pm 1, 0, 0)$ we have $\nabla G // \mathbf{e}_1$. Hence $T_{(\pm 1,0,0)}S = \operatorname{Span}\{\mathbf{e}_2, \mathbf{e}_3\}$. Since $H\mathcal{L}(\pm 1, 0, 0) = -6I_2$ on this subspace, we conclude by Theorem 3 that these points are local maxima for f on S.

At the points $(0, \pm\frac{1}{\sqrt{2}}, 0)$ we have $\mathcal{L}(x, y, z) = f(x, y, z) + \frac{1}{2}G(x, y, z) = \frac{3}{2}x^2 + \frac{3}{2}x^2 - 1$ and

$$H\mathcal{L}(0, \pm\frac{1}{\sqrt{2}}, 0) = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 3 \end{pmatrix}.$$

Since $\nabla G(0, \pm\frac{1}{\sqrt{2}}, 0) // \mathbf{e}_2$, the tangent space at these points is $\operatorname{Span}\{\mathbf{e}_1, \mathbf{e}_3\}$. On this subspace we have $H\mathcal{L}(0, \pm\frac{1}{\sqrt{2}}, 0) = 3I_2$, so we see that these points are local minima for f on S.

Finally, consider the points $(0, 0, \pm\frac{1}{\sqrt{3}})$. Here we have $\lambda = 0$ so that $\mathcal{L}(x, y, z) = f(x, y, z) = x^2 - y^2$. The Hessian is given by

$$H\mathcal{L}(0, \pm\frac{1}{\sqrt{2}}, 0) = \begin{pmatrix} 2 & 0 & 0 \\ 0 & -2 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Similarly to the previous cases, $\nabla G(0, 0, \pm\frac{1}{\sqrt{3}}) // \mathbf{e}_3$ so that $T_{(0,0,\pm\frac{1}{\sqrt{3}})}S = \operatorname{Span}\{\mathbf{e}_1, \mathbf{e}_2\}$. However, we have $\mathbf{e}_1^T H\mathcal{L}(0, 0, \pm\frac{1}{\sqrt{3}})\mathbf{e}_1 = 2$ and $\mathbf{e}_2^T H\mathcal{L}(0, 0, \pm\frac{1}{\sqrt{3}})\mathbf{e}_2 = -2$, so Theorem 3 confirms that the points $(0, 0, \pm\frac{1}{\sqrt{3}})$ are neither local minima nor local maxima for f on S.

The sufficient conditions given in Theorem 3 can sometimes be difficult to verify in practice. Another standard approach is to associate to any constrained optimization problem a general Lagrangian function $\mathcal{L}$, and consider what is often called the *bordered Hessian* of $\mathcal{L}$. Specifically, given a function f to optimize subject to the constraints $\mathbf{G}$, define the Lagrangian $\mathcal{L} \colon \mathbb{R}^{n+m} \to \mathbb{R}$ by

$$\mathcal{L}(y_1, \ldots, y_m, x_1, \ldots, x_n) = f(x_1, \ldots, x_n) + \sum_{k=1}^{m} y_k G_k(x_1, \ldots, x_n).$$

Then by Theorem 2, if a point $\mathbf{a}$ is a solution to the optimization problem with Lagrange multiplier $\lambda = (\lambda_1, \ldots, \lambda_m)$, then $(\lambda, \mathbf{a})$ is a critical point of $\mathcal{L}$. The bordered Hessian of the problem is then the Hessian matrix of this Lagrangian function. Thus the bordered Hessian at the point $(\lambda, \mathbf{a})$ is given by

$$H\mathcal{L}(\lambda, \mathbf{a}) = \begin{pmatrix} \mathbf{0}_{m\times m} & -D\mathbf{G}(\mathbf{a}) \\ -D\mathbf{G}(\mathbf{a})^T & H_{\mathbf{x}}\mathcal{L}(\mathbf{a}) \end{pmatrix},$$

where $\mathbf{0}_{m\times m}$ denotes the $m \times m$ matrix of zeroes, and $H_{\mathbf{x}}\mathcal{L}(\mathbf{a})$ denotes the $n \times n$ matrix whose (i,j)-entry is $\partial_{x_i}\partial_{x_j}\mathcal{L}(\mathbf{a})$. If $\mathbf{a} \in S$, then the matrix $D\mathbf{G}(\mathbf{a})$ has rank m, and so it has an invertible $m \times m$ submatrix. By relabelling if necessary, we may assume that the matrix

$$\begin{pmatrix} \partial_1 G_1(\mathbf{a}) & \dots & \partial_m G_1(\mathbf{a}) \\ \vdots & \ddots & \vdots \\ \partial_1 G_n(\mathbf{a}) & \dots & \partial_m G_m(\mathbf{a}) \end{pmatrix}$$

is invertible. Then the theory of quadratic forms shows that in the case where $H\mathcal{L}(\mathbf{a})$ is invertible, parts (3)-(5) of Theorem 3 are equivalent to the following theorem.

Theorem 4. *Let $f\colon \mathbb{R}^n \to \mathbb{R}$ and $\mathbf{G}\colon \mathbb{R}^n \to \mathbb{R}^m$ be of class C^2 and denote by $\mathcal{L}$ the Lagrangian function of the corresponding optimization problem. Assume $\mathbf{0}$ is a regular value for $\mathbf{G}$, and set $S = \mathbf{G}^{-1}(\mathbf{0})$. Suppose $(\lambda, \mathbf{a}) \in \mathbb{R}^{n+m}$ is a critical point of $\mathcal{L}$ and that the matrix*

$$\begin{pmatrix} \partial_1 G_1(\mathbf{a}) & \dots & \partial_m G_1(\mathbf{a}) \\ \vdots & \ddots & \vdots \\ \partial_1 G_n(\mathbf{a}) & \dots & \partial_m G_m(\mathbf{a}) \end{pmatrix}$$

is invertible. Let $H\mathcal{L}(\lambda, \mathbf{a})$ be the bordered Hessian of $\mathcal{L}$ at $(\lambda, \mathbf{a})$, and suppose $\det H\mathcal{L}(\lambda, \mathbf{a}) \neq 0$. *Denote by H_j the upper leftmost $j \times j$ submatrix of $H\mathcal{L}(\lambda, \mathbf{a})$, and set $d_j = \det H_j$. Consider the sequence of $n - m$ numbers*

$$(-1)^m d_{2m+1},\ (-1)^m d_{2m+2},\ \dots,\ (-1)^m d_{m+n}. \tag{3}$$

(1) *If the sequence in* (3) *consists entirely of positive numbers, then f has a local minimum on S at $\mathbf{a}$.*
(2) *If the sequence in* (3) *begins with a negative number and thereafter alternates in sign, then f has a local maximum on S at $\mathbf{a}$.*
(3) *If neither of the above conditions holds, then f has neither a local minimum nor a local maximum on S at $\mathbf{a}$.*

We would like to end these notes with a beautiful application of Lagrange multipliers.

Example 5. In our notes on unconstrained optimization, we made use of a very important theorem from linear algebra known as the Spectral Theorem, but we did not provide a proof. Most proofs of this theorem rely on complex numbers and the use of the Fundamental Theorem of Algebra. Using the method of Lagrange multipliers, we can provide a simple proof of this theorem in the case of real, symmetric matrices. Recall that the theorem states that if A is a real symmetric $n \times n$ matrix, then there exists an orthonormal basis of $\mathbb{R}^n$ consisting of eigenvectors of A.

To prove this theorem, consider the function $f\colon \mathbb{R}^n \to \mathbb{R}$ defined by $f(\mathbf{x}) = \mathbf{x}^T A\mathbf{x}$. We will solve the problem of optimizing f on the unit sphere $S^{n-1} \subseteq \mathbb{R}^n$. Since S^{n-1} is compact, we know that f achieves a minimum somewhere on the sphere, say $\mathbf{v}_1$ minimizes f on S^{n-1}. If we set $G(\mathbf{x}) = \sum_{i=1}^n x_i^2 - 1$ then $S^{n-1} = G^{-1}(0)$ and we have $\nabla f(\mathbf{x}) = 2A\mathbf{x}$ and $\nabla G(\mathbf{x}) = 2\mathbf{x}$ (which also shows that 0 is a regular value of G). It follows from Theorem 2 that for some $\lambda_1 \in \mathbb{R}$ we have

$$A\mathbf{v}_1 = \lambda_1 \mathbf{v}_1.$$

In other words, $\mathbf{v}_1$ is an eigenvector of A. Set $W = (\text{Span}(\mathbf{v}_1))^\perp$. If $\mathbf{w} \in W$, then by symmetry of A we have

$$\mathbf{v}_1^T(A\mathbf{w}) = (\mathbf{v}_1^T A^T)\mathbf{w} = \lambda_1 \mathbf{v}_1^T \mathbf{w} = 0,$$

so that $A\mathbf{w} \in W$. Thus we can restrict A to W and obtain a real $(n-1) \times (n-1)$ symmetric matrix. Repeating the argument yields an eigenvector $\mathbf{v}_2 \in W \cap S^{n-2}$, and we can show that $A|_W$ preserves the subspace $(\text{Span}(\mathbf{v}_2))^\perp \subseteq W$. Continuing in this way we obtain a set n orthonormal eigenvectors of A, completing the proof of the theorem.

References

[1] Pressley, Andrew, *Elementary Differential Geometry*, Springer Undergraduate Mathematics Series, Springer-Verlag London Ltd., London, 2010. MR2598317.