

Solving Equality-Constrained Optimization Problems via Perturbation Analysis

Su Jianlin

October 3, 2026

For smooth functions without constraints, the universal approach to finding extrema is of course to take the derivative (gradient) and set it equal to zero. But what about extremum problems with equality constraints? Many readers should be familiar with this as well—the standard method is the Lagrange multiplier method. However, Lagrange multipliers always feel like a purely algebraic trick, lacking the intuitive flavor of perturbation analysis that one has when finding unconstrained extrema.

In this article, we give it a try: if we carry the perturbation idea for unconstrained extrema over into constrained optimization, what result do we get?

1 Stationarity Condition

First, let us review why the gradient being zero is a necessary condition for an unconstrained extremum (the stationarity condition). Consider the objective function $f(\mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^d$, and suppose we want to find the minimizer of $f(\mathbf{x})$. We make a first-order approximation near $\mathbf{x}$:

$$f(\mathbf{x} + \delta\mathbf{x}) = f(\mathbf{x}) + \delta\mathbf{x}^\top \nabla_{\mathbf{x}} f(\mathbf{x}) \quad (1)$$

If $f(\mathbf{x})$ is the minimum, then for any perturbation $\delta\mathbf{x} \neq \mathbf{0}$, it must always hold that $f(\mathbf{x} + \delta\mathbf{x}) \geq f(\mathbf{x})$, which means $\delta\mathbf{x}^\top \nabla_{\mathbf{x}} f(\mathbf{x}) \geq 0$. But $\delta\mathbf{x}^\top \nabla_{\mathbf{x}} f(\mathbf{x})$ is a linear function of $\delta\mathbf{x}$; for it to be always non-negative, the only possibility is

$$\nabla_{\mathbf{x}} f(\mathbf{x}) = \mathbf{0} \quad (2)$$

The same discussion applies to maxima, so this is the stationarity condition shared by maxima and minima. When the stationarity condition holds, the next leading term of the approximate expansion is

$$f(\mathbf{x} + \delta\mathbf{x}) = f(\mathbf{x}) + \frac{1}{2} \delta\mathbf{x}^\top \nabla_{\mathbf{x}}^2 f(\mathbf{x}) \delta\mathbf{x} \quad (3)$$

By the same principle, for $f(\mathbf{x} + \delta\mathbf{x}) \geq f(\mathbf{x})$ to always hold, we need $\delta\mathbf{x}^\top \nabla_{\mathbf{x}}^2 f(\mathbf{x}) \delta\mathbf{x} \geq 0$ to always hold, i.e., the Hessian matrix $\nabla_{\mathbf{x}}^2 f(\mathbf{x})$ is positive semi-definite, which is a necessary condition for a minimum; conversely, the necessary condition for a maximum is that $\nabla_{\mathbf{x}}^2 f(\mathbf{x})$ is negative semi-definite.

Of course, this discussion is not a rigorous mathematical proof; it can only be regarded as a heuristic derivation.

2 Equality Constraints

Now let us see what changes in the perturbation process when equality constraints are introduced. Consider

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{g}(\mathbf{x}) = \mathbf{0} \quad (4)$$

where $\mathbf{g}(\mathbf{x}) \in \mathbb{R}^k$ represents k equality constraints. Pick any $\mathbf{x}$ satisfying $\mathbf{g}(\mathbf{x}) = \mathbf{0}$; similarly, considering the first-order approximation of the perturbation gives

$$f(\mathbf{x} + \delta\mathbf{x}) = f(\mathbf{x}) + \delta\mathbf{x}^\top \nabla_{\mathbf{x}} f(\mathbf{x}) \quad (5)$$

Similarly, the necessary condition for $f(\mathbf{x})$ to be the minimum is still that $\delta\mathbf{x}^\top \nabla_{\mathbf{x}} f(\mathbf{x}) \geq 0$ always holds, but the difference is that now it does not need to hold for arbitrary perturbations $\delta\mathbf{x}$, but only for all $\delta\mathbf{x}$ satisfying $\mathbf{g}(\mathbf{x} + \delta\mathbf{x}) = \mathbf{0}$. Making a first-order approximation of $\mathbf{g}(\mathbf{x} + \delta\mathbf{x}) = \mathbf{0}$ gives

$$\mathbf{J}\delta\mathbf{x} = \mathbf{0}, \quad \mathbf{J} = \nabla_{\mathbf{x}} \mathbf{g}(\mathbf{x}) \in \mathbb{R}^{k \times d} \quad (6)$$

Note that $\mathbf{J}\delta\mathbf{x} = \mathbf{0}$ is a linear system of equations in $\delta\mathbf{x}$, and its general solution can be written using the “pseudoinverse”:

$$\delta\mathbf{x} = (\mathbf{I} - \mathbf{J}^\dagger \mathbf{J})\delta\mathbf{z} \quad (7)$$

where $\mathbf{J}^\dagger$ is the pseudoinverse of $\mathbf{J}$, and $\delta\mathbf{z}$ denotes an arbitrary unconstrained perturbation vector. When $\mathbf{J}\mathbf{J}^\top$ is invertible, $\mathbf{J}^\dagger = \mathbf{J}^\top(\mathbf{J}\mathbf{J}^\top)^{-1}$, from which we see that $\mathbf{J}^\dagger \mathbf{J}$ is a symmetric matrix. Substituting this into $\delta\mathbf{x}^\top \nabla_{\mathbf{x}} f(\mathbf{x}) \geq 0$ and using the arbitrariness of $\delta\mathbf{z}$, we obtain the stationarity condition

$$(\mathbf{I} - \mathbf{J}^\dagger \mathbf{J})\nabla_{\mathbf{x}} f(\mathbf{x}) = \mathbf{0} \quad (8)$$

Here $(\mathbf{I} - \mathbf{J}^\dagger \mathbf{J})\nabla_{\mathbf{x}} f(\mathbf{x})$ is called the projected gradient of $f(\mathbf{x})$ on the manifold $\mathbf{g}(\mathbf{x}) = \mathbf{0}$; it is the orthogonal projection of the Euclidean gradient $\nabla_{\mathbf{x}} f(\mathbf{x})$ onto the tangent space. In the language of Riemannian geometry, it is also called the Riemannian gradient.

3 Lagrange

Now let us compare this with the Lagrange multiplier method. The Lagrange multiplier method introduces undetermined coefficients $\mathbf{c} \in \mathbb{R}^k$ and changes the objective function to

$$F(\mathbf{x}, \mathbf{c}) = f(\mathbf{x}) - \mathbf{c}^\top \mathbf{g}(\mathbf{x}) \quad (9)$$

and then seeks the stationary points of the unconstrained objective $F(\mathbf{x}, \mathbf{c})$. Setting the gradients on both sides to zero directly gives

$$\nabla_{\mathbf{x}} F(\mathbf{x}, \mathbf{c}) = \nabla_{\mathbf{x}} f(\mathbf{x}) - \mathbf{J}^\top \mathbf{c} = \mathbf{0}, \quad \nabla_{\mathbf{c}} F(\mathbf{x}, \mathbf{c}) = -\mathbf{g}(\mathbf{x}) = \mathbf{0} \quad (10)$$

Note that the pseudoinverse satisfies $\mathbf{J}^\dagger \mathbf{J} = \mathbf{J}^\top(\mathbf{J}^\top \mathbf{J})^\dagger$, so Equation (8) can also be written as $\nabla_{\mathbf{x}} f(\mathbf{x}) - \mathbf{J}^\top(\mathbf{J}^\top \mathbf{J})^\dagger \nabla_{\mathbf{x}} f(\mathbf{x}) = \mathbf{0}$. If we denote $(\mathbf{J}^\top \mathbf{J})^\dagger \nabla_{\mathbf{x}} f(\mathbf{x})$ as $\mathbf{c}$, then it is exactly the same as the above, i.e., from Equation (8) we can derive the same condition as the Lagrange multiplier method.

Conversely, can we derive Equation (8) from the Lagrange multiplier method? Naturally, yes. Assume $k < d$ and that $\mathbf{J}$ has rank k ; then regarding $\nabla_{\mathbf{x}} f(\mathbf{x}) - \mathbf{J}^\top \mathbf{c} = \mathbf{0}$ as an equation in $\mathbf{c}$, it is an overdetermined system, and if it has a solution it can only be $\mathbf{c} = (\mathbf{J}^\top \mathbf{J})^\dagger \nabla_{\mathbf{x}} f(\mathbf{x})$; substituting back into the equation yields Equation (8).

This shows that the conditions given by the two are completely equivalent mathematically. Of course, the Lagrange multiplier formulation looks cleaner, but the author believes that the perturbation approach is more consistent with the unconstrained case, and therefore prefers this way of understanding it.

Moreover, the simplicity of the Lagrange multiplier method is only formal; to truly understand it is no easier than the perturbation perspective. Readers can entirely regard the process from the perturbation idea to the projected gradient as a derivation of the Lagrange multiplier method.

4 Two Examples

We use the following two examples to deepen our understanding.

4.1 Sphere Constraint

As the first example, consider the simple extremum on a sphere:

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{s.t.} \quad \|\mathbf{x}\|_2 = 1 \quad (11)$$

It is easy to compute $\nabla_{\mathbf{x}} \|\mathbf{x}\|_2 = \mathbf{x} / \|\mathbf{x}\|_2 = \mathbf{x}$, so $\mathbf{J} = \mathbf{x}^\top$; since $\mathbf{x}$ is a unit vector, $\mathbf{J}^\dagger = \mathbf{x}$. Therefore the stationarity condition is

$$\mathbf{0} = (\mathbf{I} - \mathbf{x}\mathbf{x}^\top) \nabla_{\mathbf{x}} f(\mathbf{x}) = \nabla_{\mathbf{x}} f(\mathbf{x}) - \mathbf{x}(\mathbf{x}^\top \nabla_{\mathbf{x}} f(\mathbf{x})) \quad (12)$$

The final expression is precisely the projection of the gradient onto the tangent space of the sphere.

4.2 Orthogonality Constraint

Another example is the optimization problem on the Stiefel manifold:

$$\min_{\mathbf{X}} f(\mathbf{X}) \quad \text{s.t.} \quad \mathbf{X}^\top \mathbf{X} = \mathbf{I}_m, \quad \mathbf{X} \in \mathbb{R}^{n \times m} (n \geq m) \quad (13)$$

If one only wants the stationarity condition, using the Lagrange multiplier method would be much faster, but here we still work out the projected gradient in full. The difficulty here is that $\mathbf{X}$ itself is a matrix, so $\mathbf{J}$ would be a higher-order tensor, which imposes a considerable mental burden.

To deal with this, we adopt the linear-operator perspective: making a first-order approximation of $\mathbf{X}^\top \mathbf{X} - \mathbf{I}_m$ gives $\delta \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \delta \mathbf{X}$, which according to the previous derivation should equal $\mathbf{J} \delta \mathbf{x}$. We directly define the linear operator

$$\mathbf{J}[\mathbf{V}] = \mathbf{V}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{V} \quad (14)$$

Next we need to derive $\mathbf{J}^\top[\mathbf{V}]$, understood via $\langle \mathbf{U}, \mathbf{J}^\top[\mathbf{V}] \rangle_F = \langle \mathbf{V}, \mathbf{J}[\mathbf{U}] \rangle_F$ (analogous to $\mathbf{u}^\top \mathbf{J}^\top \mathbf{v} = \mathbf{v}^\top \mathbf{J} \mathbf{u}$). We already know $\mathbf{J}[\mathbf{U}]$ and $\langle \cdot, \cdot \rangle_F$, so

$$\langle \mathbf{U}, \mathbf{J}^\top[\mathbf{V}] \rangle_F = \text{tr}(\mathbf{V}^\top (\mathbf{U}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{U})) = \text{tr}(\mathbf{U}^\top \mathbf{X} (\mathbf{V} + \mathbf{V}^\top)) = \langle \mathbf{U}, \mathbf{X}(\mathbf{V} + \mathbf{V}^\top) \rangle_F \quad (15)$$

From this we can read off $\mathbf{J}^\top[\mathbf{V}] = \mathbf{X}(\mathbf{V} + \mathbf{V}^\top)$. According to $\mathbf{J}^\dagger \mathbf{J} = \mathbf{J}^\top (\mathbf{J} \mathbf{J}^\top)^{-1} \mathbf{J}$, we also need $\mathbf{J} \mathbf{J}^\top$ and $(\mathbf{J} \mathbf{J}^\top)^{-1}$:

$$\mathbf{J} \mathbf{J}^\top[\mathbf{V}] = \mathbf{J}[\mathbf{X}(\mathbf{V} + \mathbf{V}^\top)] = (\mathbf{V} + \mathbf{V}^\top) \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} (\mathbf{V} + \mathbf{V}^\top) = 2(\mathbf{V} + \mathbf{V}^\top) \quad (16)$$

If $\mathbf{V}$ is a symmetric matrix, then $\mathbf{J}\mathbf{J}^\top[\mathbf{V}] = 4\mathbf{V}$, so $(\mathbf{J}\mathbf{J}^\top)^{-1}[\mathbf{V}] = \mathbf{V}/4$; and $\mathbf{J}[\mathbf{V}]$ happens to be symmetric, so we only need this special case. Putting it all together,

$$\mathbf{J}^\dagger \mathbf{J}[\mathbf{V}] = \mathbf{J}^\top (\mathbf{J}\mathbf{J}^\top)^{-1} \mathbf{J}[\mathbf{V}] = \mathbf{J}^\top [(\mathbf{V}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{V})/4] = \mathbf{X}(\mathbf{V}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{V})/2 \quad (17)$$

Let $\mathbf{G} = \nabla_{\mathbf{X}} f(\mathbf{X})$; then its projected gradient on the Stiefel manifold is

$$\mathbf{G} - \mathbf{J}^\dagger \mathbf{J}[\mathbf{G}] = \mathbf{G} - \mathbf{X}[\mathbf{X}^\top \mathbf{G}]_{\text{sym}} \quad (18)$$

where $[\mathbf{Z}]_{\text{sym}} = (\mathbf{Z} + \mathbf{Z}^\top)/2$ is the symmetrization operator. Setting this equal to zero and left-multiplying both sides by $\mathbf{X}^\top$ gives $\mathbf{X}^\top \mathbf{G} = [\mathbf{X}^\top \mathbf{G}]_{\text{sym}}$; that is to say, the stationarity condition is that $\mathbf{X}^\top \mathbf{G}$ is itself a symmetric matrix!

5 Second-Order Conditions

So, under equality-constrained optimization, is there a method analogous to using the positive/negative definiteness of $\nabla_{\mathbf{x}}^2 f(\mathbf{x})$ to determine minima/maxima? Naturally there is, and there is a rather elegant derivation. As with the Lagrange multiplier method, we define

$$F(\mathbf{x}, \mathbf{c}) = f(\mathbf{x}) - \mathbf{c}^\top \mathbf{g}(\mathbf{x}) \quad (19)$$

For any stationary point $\mathbf{x}$, we automatically have $\mathbf{g}(\mathbf{x}) = \mathbf{0}$, so for any $\mathbf{c} \in \mathbb{R}^k$ we have $F(\mathbf{x}, \mathbf{c}) \equiv f(\mathbf{x})$. Thus, expanding to second order,

$$f(\mathbf{x} + \delta \mathbf{x}) = F(\mathbf{x} + \delta \mathbf{x}, \mathbf{c}) = \underbrace{F(\mathbf{x}, \mathbf{c})}_{\equiv f(\mathbf{x})} + \delta \mathbf{x}^\top (\nabla_{\mathbf{x}} f(\mathbf{x}) - \mathbf{J}^\top \mathbf{c}) + \frac{1}{2} \delta \mathbf{x}^\top \mathbf{H}(\mathbf{c}) \delta \mathbf{x} \quad (20)$$

where $\mathbf{H}[\mathbf{c}] \triangleq \nabla_{\mathbf{x}}^2 (f(\mathbf{x}) - \mathbf{c}^\top \mathbf{g}(\mathbf{x}))$. Due to the arbitrariness of $\mathbf{c}$, we can simply take $\mathbf{c} = (\mathbf{J}^\top)^\dagger \nabla_{\mathbf{x}} f(\mathbf{x})$ so that the first-order term vanishes, giving

$$f(\mathbf{x} + \delta \mathbf{x}) = f(\mathbf{x}) + \frac{1}{2} \delta \mathbf{x}^\top \mathbf{H}[(\mathbf{J}^\top)^\dagger \nabla_{\mathbf{x}} f(\mathbf{x})] \delta \mathbf{x} \quad (21)$$

Now only the second-order term remains, so the approximation accuracy for $\delta \mathbf{x}$ only needs to be first order, i.e., $\delta \mathbf{x} = (\mathbf{I} - \mathbf{J}^\dagger \mathbf{J}) \delta \mathbf{z}$ suffices. Substituting this in gives

$$f(\mathbf{x} + \delta \mathbf{x}) = f(\mathbf{x}) + \frac{1}{2} \delta \mathbf{z}^\top (\mathbf{I} - \mathbf{J}^\dagger \mathbf{J}) \mathbf{H}[(\mathbf{J}^\top)^\dagger \nabla_{\mathbf{x}} f(\mathbf{x})] (\mathbf{I} - \mathbf{J}^\dagger \mathbf{J}) \delta \mathbf{z} \quad (22)$$

By the arbitrariness of $\delta \mathbf{z}$, we know that a necessary condition for $f(\mathbf{x})$ to be a minimum is that $(\mathbf{I} - \mathbf{J}^\dagger \mathbf{J}) \mathbf{H}[(\mathbf{J}^\top)^\dagger \nabla_{\mathbf{x}} f(\mathbf{x})] (\mathbf{I} - \mathbf{J}^\dagger \mathbf{J})$ is positive semi-definite. Similarly, for a maximum this matrix must be negative semi-definite.

Of course, whether for unconstrained or constrained optimization, judging maxima/minima via the positive or negative definiteness of the Hessian matrix usually has only theoretical value; in practice one typically filters out all locations satisfying the stationarity condition and then analyzes each case on its own merits.

6 Summary

In this article, we applied the perturbation analysis idea for unconstrained extrema to equality-constrained optimization, derived the projected gradient on the tangent space, and found that the stationarity condition for extrema is that the corresponding projected gradient is zero. Finally, we also proved its equivalence with the Lagrange multiplier method in terms of conclusions.

When reprinting, please include the address of this article:

<https://kexue.fm/archives/11928>

For more detailed reprinting matters, please refer to: “Scientific Spaces FAQ”