

Review of High-Dimensional Calculus

Ju Sun*

January 25, 2026

High-dimensional calculus is typically not covered in basic calculus courses, but it is the language of modern machine learning—we almost always express quantities of interest as vectors, matrices, or even tensors. In this set of notes, we quickly go over the basics of high-dimensional calculus that are most useful for machine learning and highlight certain computational techniques that are not often taught elsewhere. Two recommended sources for learning high-dimensional calculus are [Mun97, Col12]. [Zor15, Zor16] is another set of useful references. There is also an interesting online tool for computing symbolic matrix derivatives <https://www.matrixcalculus.org/>.

1 Our notations

- scalars: x (small letters) vectors: \mathbf{x} (bold small) matrices: \mathbf{X} (bold capital) tensors: \mathcal{X} (script capital) sets: S (capital)
- vectors are always **column vectors**, unless stated otherwise
- x_i : i -th element of \mathbf{x} x_{ij} : (i, j) -th element of \mathbf{X} \mathbf{x}^i : i -th row of \mathbf{X} as a **row vector** x_j : j -th column of \mathbf{X} as a **column vector**
- \mathbb{R} : real numbers \mathbb{R}_+ : positive reals \mathbb{R}^n : space of n -dimensional vectors $\mathbb{R}^{m \times n}$: space of $m \times n$ matrices $\mathbb{R}^{m \times n \times k}$: space of $m \times n \times k$ tensors, etc
- $[n] \doteq \{1, \dots, n\}$ (often used by theoretical computer scientists; \doteq means “defined as”)
- $[v_i]_i$ means a vector \mathbf{v} whose i -th entry takes the value v_i , e.g., $[\frac{1}{i^2}]_i = [1; \frac{1}{4}; \frac{1}{9}; \dots; \frac{1}{100}]$ if the vector length is 10; $[m_{ij}]_{i,j}$ means a matrix \mathbf{M} whose (i, j) -th entry takes the value m_{ij} .

2 Differentiability

2.1 First-order differentiability

Definition 2.1 (First-order derivative or Jacobian). *Consider a function $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^m$. f is (Fréchet) differentiable at a point \mathbf{x}_0 if there exists a matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$ such that*

$$\lim_{\delta \rightarrow 0} \frac{f(\mathbf{x}_0 + \boldsymbol{\delta}) - f(\mathbf{x}_0) - \mathbf{B}\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|_2} = \mathbf{0}, \quad (2.1)$$

or equivalently,

$$f(\mathbf{x}_0 + \boldsymbol{\delta}) = f(\mathbf{x}_0) + \mathbf{B}\boldsymbol{\delta} + \underbrace{o(\|\boldsymbol{\delta}\|_2)}_{\text{a residual term lower-order than } \boldsymbol{\delta}} \quad \text{as } \boldsymbol{\delta} \rightarrow \mathbf{0}. \quad (2.2)$$

Here \mathbf{B} is called the (Fréchet) derivative, or the Jacobian of f at \mathbf{x}_0 , denoted as $\mathbf{J}_f(\mathbf{x}_0) \in \mathbb{R}^{m \times n}$.

*Department of Computer Science and Engineering, University of Minnesota at Twin Cities. Email: jusun@umn.edu.

The equivalence inside the theorem is due to the following fact: for any vector-valued function $h(\boldsymbol{\delta}) : \mathbb{R}^n \rightarrow \mathbb{R}^m$,

$$h(\boldsymbol{\delta}) \in o(\|\boldsymbol{\delta}\|_2) \text{ as } \boldsymbol{\delta} \rightarrow \mathbf{0} \iff \lim_{\boldsymbol{\delta} \rightarrow \mathbf{0}} \frac{h(\boldsymbol{\delta})}{\|\boldsymbol{\delta}\|_2} = \mathbf{0} \iff \lim_{\boldsymbol{\delta} \rightarrow \mathbf{0}} \frac{\|h(\boldsymbol{\delta})\|_2}{\|\boldsymbol{\delta}\|_2} = 0. \quad (2.3)$$

If $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable at a point \mathbf{x}_0 , all first-order partial derivatives $\frac{\partial f_i}{\partial x_j}(\mathbf{x}_0)$ for all i, j exist. We also have

$$\mathbf{J}_f(\mathbf{x}_0) = \left[\frac{\partial f_i}{\partial x_j}(\mathbf{x}_0) \right]_{i,j}, \quad (2.4)$$

i.e., $\mathbf{J}_f(\mathbf{x}_0)$ is the collection of all first-order partial derivatives $\frac{\partial f_i}{\partial x_j}(\mathbf{x}_0)$ for all i, j .

Definition 2.2 (Gradient). Consider a function $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$, i.e., scalar-valued. The gradient $\nabla f(\mathbf{x}) \in \mathbb{R}^{n \times 1}$, which is a column vector, is the transpose of the Jacobian $\mathbf{J}_f(\mathbf{x}) \in \mathbb{R}^{1 \times n}$.

According to this convention, for $f(\mathbf{X}) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$, the gradient should be $\mathbb{R}^{(m \times n) \times 1}$, which is a length- $(m \times n)$ vector. This is *inconvenient* for many purposes. So in practice, this vector is reshaped into an $m \times n$ matrix, i.e., the same shape as that of \mathbf{X} . We will follow this convention, i.e., $\nabla f(\mathbf{X}) = \left[\frac{\partial f}{\partial x_{i,j}}(\mathbf{X}) \right]_{i,j} \in \mathbb{R}^{m \times n}$.

Theorem 2.3 (Calculus rules of Jacobian). Assume $f(\mathbf{x}), g(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are differentiable at a point $\mathbf{x}_0 \in \mathbb{R}^n$.

- **linearity:** for any $\lambda_1, \lambda_2 \in \mathbb{R}$, $\lambda_1 f + \lambda_2 g$ is differentiable at \mathbf{x}_0 and $\mathbf{J}_{\lambda_1 f + \lambda_2 g}(\mathbf{x}_0) = \lambda_1 \mathbf{J}_f(\mathbf{x}_0) + \lambda_2 \mathbf{J}_g(\mathbf{x}_0)$;
- **product:** assume $m = 1$, fg is differentiable at \mathbf{x}_0 and $\mathbf{J}_{fg}(\mathbf{x}_0) = f(\mathbf{x}_0) \mathbf{J}_g(\mathbf{x}_0) + g(\mathbf{x}_0) \mathbf{J}_f(\mathbf{x}_0)$;
- **quotient:** assume $m = 1$ and $g(\mathbf{x}_0) \neq 0$, $\frac{f}{g}$ is differentiable at \mathbf{x}_0 and $\mathbf{J}_{\frac{f}{g}}(\mathbf{x}_0) = \frac{g(\mathbf{x}_0) \mathbf{J}_f(\mathbf{x}_0) - f(\mathbf{x}_0) \mathbf{J}_g(\mathbf{x}_0)}{g^2(\mathbf{x}_0)}$;
- **Chain rule:** Let $f(\mathbf{x}) : \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $h(\mathbf{y}) : \mathbb{R}^n \rightarrow \mathbb{R}^k$. If f is differentiable at \mathbf{x}_0 , and h is differentiable at \mathbf{y}_0 where $\mathbf{y}_0 = f(\mathbf{x}_0)$. Then, $h \circ f(\mathbf{x}) : \mathbb{R}^m \rightarrow \mathbb{R}^k$ is differentiable at \mathbf{x}_0 , and

$$\mathbf{J}_{h \circ f}(\mathbf{x}_0) = \mathbf{J}_h(f(\mathbf{x}_0)) \mathbf{J}_f(\mathbf{x}_0). \quad (2.5)$$

When $k = 1$,

$$\nabla[h \circ f](\mathbf{x}_0) = \mathbf{J}_f^\top(\mathbf{x}_0) \nabla h(f(\mathbf{x}_0)). \quad (2.6)$$

2.2 Deriving the Jacobian

There are two main methods. One is invoking Eq. (2.4) and the calculus rules in Theorem 2.3. The other is by the perturbation-expansion method based on Definition 2.1, as explained below. Let's take an example $f(\mathbf{x}) = \mathbf{A}\mathbf{x} - \mathbf{b}$ for $\mathbf{A} \in \mathbb{R}^{n \times m}$ and $\mathbf{b} \in \mathbb{R}^n$ and try to derive the Jacobian.

- **Elementwise calculation and calculus rules.** Obviously $\mathbf{J}_b(\mathbf{x}) = \mathbf{0}$ so we can focus on the $\mathbf{A}\mathbf{x}$ term. Now

$$\frac{\partial(\mathbf{A}\mathbf{x})_i}{\partial x_j} = \frac{\partial(\mathbf{a}^i \mathbf{x})}{\partial x_j} = \frac{\partial \sum_k a_{ik} x_k}{\partial x_j} = a_{ij}. \quad (2.7)$$

So $\mathbf{J}_{\mathbf{A}\mathbf{x}}(\mathbf{x}) = [a_{ij}]_{i,j} = \mathbf{A}$, and by linearity $\mathbf{J}_f(\mathbf{x}) = \mathbf{A}$.

- **Perturbation-expansion method.** We make an infinitesimal perturbation δ to x , so

$$f(x + \delta) = A(x + \delta) - b = (Ax - b) + A\delta = f(x) + A\delta. \quad (2.8)$$

Comparing this to [Eq. \(2.2\)](#) in [Definition 2.1](#), we easily obtain that $J_f(x) = A$. Note that in the perturbation-expansion method, after the infinitesimal perturbation δ , we rearrange the terms to match the form of [Eq. \(2.2\)](#), i.e.,

$$f(x) + \text{linear term in } \delta + \text{lower-order term in } \|\delta\|_2 \quad (2.9)$$

so that we can read off the Jacobian from the linear term.

Example 2.4. Consider the least-squares objective $f(x) = \|Ax - b\|_2^2$. We will derive the Jacobian, which is the transpose of the gradient.

- **Chain rule.** We can view f as composition of $g(x) = Ax - b$ and $h(y) = \|y\|_2^2$ so that $f = h \circ g(x)$. From our last example, $J_g(x) = A$. For h , it is easy to check that $J_h(y) = 2y^\top$. Applying the chain rule, we obtain that

$$J_f(x) = J_h(Ax - b)J_g(x) = 2(Ax - b)^\top A. \quad (2.10)$$

- **Perturbation-expansion method.** By making an infinitesimal perturbation δ to x , we obtain

$$f(x + \delta) = \|A(x + \delta) - b\|_2^2 \quad (2.11)$$

$$= \|(Ax - b) + A\delta\|_2^2 \quad (2.12)$$

$$= \|Ax - b\|_2^2 + \|A\delta\|_2^2 + 2\langle Ax - b, A\delta \rangle. \quad (2.13)$$

Let us clarify a few points before proceeding. We use $\langle \cdot, \cdot \rangle$ to mean the inner product for vectors, i.e., for $u, v \in \mathbb{R}^n$, $\langle u, v \rangle \doteq \sum_i u_i v_i$. For any $p \geq 1$, the ℓ_p norm of the vector $u \in \mathbb{R}^n$ is defined as $\|u\|_p \doteq (\sum_i |u_i|^p)^{1/p}$. For $p = 2$, the norm is also called the Euclidean norm and it can be easily verified that $\|u\|_2 = \sqrt{\langle u, u \rangle}$. So for $u, v \in \mathbb{R}^n$,

$$\|u + v\|_2^2 = \langle u + v, u + v \rangle = \langle u, u \rangle + \langle v, v \rangle + 2\langle u, v \rangle = \|u\|_2^2 + \|v\|_2^2 + 2\langle u, v \rangle. \quad (2.14)$$

We have used this identity to arrive at [Eq. \(2.13\)](#). In [Eq. \(2.13\)](#), $\|Ax - b\|_2^2 = f(x)$, and $\|A\delta\|_2^2 \in O(\|\delta\|_2^2) \implies \|A\delta\|_2^2 \in o(\|\delta\|_2)$ which we do not care. The linear term is $2\langle Ax - b, A\delta \rangle$. We now invoke another identity $\langle u, v \rangle = u^\top v$ to obtain that

$$2\langle Ax - b, A\delta \rangle = 2(Ax - b)^\top A\delta. \quad (2.15)$$

Comparing this with [Eq. \(2.2\)](#), we conclude that

$$J_f(x) = 2(Ax - b)^\top A. \quad (2.16)$$

Deriving Jacobians using the chain rule is fine for simple compositions. But it quickly leads to fatigue when there are many compositions. Moreover, when intermediate variables involve matrices, intermediate Jacobians as tensors will often be involved. An example is when deriving gradients for functions involving deep neural networks, e.g.,

$$f(W) = \sum_i \|y_i - W_k \sigma(W_{k-1} \dots \sigma(W_2 \sigma(W_1 x_i)))\|_2^2. \quad (2.17)$$

So, the perturbation-expansion method becomes handy when we deal with matrix/vector variables. They allow us to derive Jacobian, gradient, and Hessian directly in matrix form, *without taking partial derivatives*.

Example 2.5 (Chain rule follows from the perturbation-expansion method). *The chain rule in Theorem 2.3 can be easily derived from the perturbation-expansion method. Consider an infinitesimal perturbation δ to \mathbf{x} in $h \circ f$:*

$$h \circ f(\mathbf{x}_0 + \delta) = h(f(\mathbf{x}_0 + \delta)) = h(f(\mathbf{x}_0) + \mathbf{J}_f(\mathbf{x}_0)\delta + o(\|\delta\|_2)), \quad (2.18)$$

where we expand $f(\mathbf{x}_0 + \delta)$ by the definition of first-order differentiability, because f is differentiable at \mathbf{x}_0 . Now $h(f(\mathbf{x}_0) + \mathbf{J}_f(\mathbf{x}_0)\delta + o(\|\delta\|_2))$ is h at the point $f(\mathbf{x}_0)$ perturbed by the infinitesimal quantity $\mathbf{J}_f(\mathbf{x}_0)\delta + o(\|\delta\|_2)$. Since h is differentiable at the point $f(\mathbf{x}_0)$, we can invoke Eq. (2.2) again and obtain that

$$\begin{aligned} & h(f(\mathbf{x}_0) + \mathbf{J}_f(\mathbf{x}_0)\delta + o(\|\delta\|_2)) \\ &= h(f(\mathbf{x}_0)) + \mathbf{J}_h(f(\mathbf{x}_0))(\mathbf{J}_f(\mathbf{x}_0)\delta + o(\|\delta\|_2)) + \underbrace{o(\|\mathbf{J}_f(\mathbf{x}_0)\delta + o(\|\delta\|_2)\|_2)}_{o(\|\delta\|_2)} \end{aligned} \quad (2.19)$$

$$= h(f(\mathbf{x}_0)) + \mathbf{J}_h(f(\mathbf{x}_0))\mathbf{J}_f(\mathbf{x}_0)\delta + \underbrace{\mathbf{J}_h(f(\mathbf{x}_0))o(\|\delta\|_2)}_{o(\|\delta\|_2)} + o(\|\delta\|_2) \quad (2.20)$$

$$= h(f(\mathbf{x}_0)) + \mathbf{J}_h(f(\mathbf{x}_0))\mathbf{J}_f(\mathbf{x}_0)\delta + o(\|\delta\|_2). \quad (2.21)$$

So $h \circ f$ is differentiable at \mathbf{x}_0 , and $\mathbf{J}_{h \circ f}(\mathbf{x}_0) = \mathbf{J}_h(f(\mathbf{x}_0))\mathbf{J}_f(\mathbf{x}_0)$.

In the above derivation, we have used several basic properties of the small-o notation:

Lemma 2.6. As $\delta \rightarrow \mathbf{0}$, (1) $o(\|f(\delta)\|) \subset o(\|h(\delta)\|)$ if $\|f(\delta)\| \leq \|h(\delta)\|$ everywhere; (2) $o(\|\delta\|) + o(\|\delta\|) = o(\|\delta\|)$; Moreover, for any matrix \mathbf{A} , (3) $o(\|\mathbf{A}\delta\|) \in o(\|\delta\|)$; (4) $\mathbf{A}o(\|\delta\|) \in o(\|\delta\|)$.

Proof. (1) Suppose that $g(\delta) \in o(\|f(\delta)\|)$, i.e., $\lim_{\delta \rightarrow \mathbf{0}} \|g(\delta)\|/\|f(\delta)\| = 0$, then

$$\lim_{\delta \rightarrow \mathbf{0}} \|g(\delta)\|/\|h(\delta)\| \leq \lim_{\delta \rightarrow \mathbf{0}} \|g(\delta)\|/\|f(\delta)\| = 0.$$

(2) Suppose that $g_1(\delta) \in o(\|\delta\|)$ and $g_2(\delta) \in o(\|\delta\|)$, then

$$\lim_{\delta \rightarrow \mathbf{0}} \|g_1(\delta) + g_2(\delta)\|/\|\delta\| \leq \lim_{\delta \rightarrow \mathbf{0}} \|g_1(\delta)\|/\|\delta\| + \lim_{\delta \rightarrow \mathbf{0}} \|g_2(\delta)\|/\|\delta\| = 0.$$

(3) Since $\|\mathbf{A}\delta\| \leq \|\mathbf{A}\|_{\text{op}}\|\delta\|$ (here $\|\mathbf{A}\|_{\text{op}}$ denotes the operator norm of \mathbf{A} , i.e., the largest singular value of \mathbf{A}), we have $o(\|\mathbf{A}\delta\|) \in o(\|\mathbf{A}\|_{\text{op}}\|\delta\|)$ by (1). Since $\|\mathbf{A}\|_{\text{op}}$ is a constant independent of δ , we can omit it as we only worry about the order of the term with respect to δ .

(4) For any $g(\delta) \in o(\|\delta\|)$, i.e., $\lim_{\delta \rightarrow \mathbf{0}} \|g(\delta)\|/\|\delta\| = 0$, we have

$$\lim_{\delta \rightarrow \mathbf{0}} \|\mathbf{A}g(\delta)\|/\|\delta\| \leq \lim_{\delta \rightarrow \mathbf{0}} \|\mathbf{A}\|_{\text{op}}\|g(\delta)\|/\|\delta\| = \|\mathbf{A}\|_{\text{op}} \lim_{\delta \rightarrow \mathbf{0}} \|g(\delta)\|/\|\delta\| = 0,$$

implying the claimed results. ■

2.3 Second-order differentiability

It is possible to define second- or even higher-order differentiability for general $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^m$. For our purposes, it suffices to consider scalar-valued functions $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$, on which we focus exclusively here. Assume f is first-order differentiable in a small ball around \mathbf{x}_0 .

- Write $\frac{\partial^2 f}{\partial x_j \partial x_i}(\mathbf{x}_0) \doteq \left[\frac{\partial}{\partial x_j} \left(\frac{\partial f}{\partial x_i} \right) \right](\mathbf{x}_0)$ provided the right side is well defined.

- **Symmetry:** If both $\frac{\partial^2 f}{\partial x_j \partial x_i}(\mathbf{x}_0)$ and $\frac{\partial^2 f}{\partial x_i \partial x_j}(\mathbf{x}_0)$ exist and both are continuous at \mathbf{x} , then they are equal.
- **Hessian (matrix):**

$$\nabla^2 f(\mathbf{x}_0) \doteq \left[\frac{\partial^2 f}{\partial x_j \partial x_i}(\mathbf{x}_0) \right]_{j,i}.$$

$\nabla^2 f$ is symmetric due to the symmetry property above.

- **Sufficient condition:** if all $\frac{\partial^2 f}{\partial x_j \partial x_i}(\mathbf{x}_0)$ exist and are **continuous** near \mathbf{x}_0 , f is 2nd-order differentiable at \mathbf{x}_0 (the converse is not true; we omit the precise definition of 2nd-order differentiability due to its technicality).

3 Taylor's theorems

Taylor's theorems take several forms. We focus on forms useful for gradient and Hessian derivation.

Theorem 3.1 (Taylor's theorem: scalar-variable version). Consider $f(x) : \mathbb{R} \rightarrow \mathbb{R}$.

- If f is 1st-order differentiable at x_0 , then

$$f(x_0 + \delta) = \underbrace{f(x_0) + \delta f'(x_0)}_{\text{first-order Taylor expansion}} + o(|\delta|) \quad \text{as } \delta \rightarrow 0. \quad (3.1)$$

- If f is 2nd-order differentiable at x_0 , then

$$f(x_0 + \delta) = \underbrace{f(x_0) + \delta f'(x_0) + \frac{1}{2} \delta^2 f''(x_0)}_{\text{second-order Taylor expansion}} + o(|\delta|^2) \quad \text{as } \delta \rightarrow 0. \quad (3.2)$$

The result can be easily generalized to scalar-valued vector- and matrix-variable functions.

Theorem 3.2 (Taylor's theorem: vector-variable version). Consider $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$.

- If f is 1st-order differentiable at \mathbf{x}_0 , then

$$f(\mathbf{x}_0 + \boldsymbol{\delta}) = \underbrace{f(\mathbf{x}_0) + \langle \nabla f(\mathbf{x}_0), \boldsymbol{\delta} \rangle}_{\text{first-order Taylor expansion}} + o(\|\boldsymbol{\delta}\|_2) \quad \text{as } \boldsymbol{\delta} \rightarrow \mathbf{0}. \quad (3.3)$$

- If f is 2nd-order differentiable at \mathbf{x}_0 , then

$$f(\mathbf{x} + \boldsymbol{\delta}) = \underbrace{f(\mathbf{x}_0) + \langle \nabla f(\mathbf{x}_0), \boldsymbol{\delta} \rangle + \frac{1}{2} \langle \boldsymbol{\delta}, \nabla^2 f(\mathbf{x}_0) \boldsymbol{\delta} \rangle}_{\text{second-order Taylor expansion}} + o(\|\boldsymbol{\delta}\|_2^2) \quad \text{as } \boldsymbol{\delta} \rightarrow \mathbf{0}. \quad (3.4)$$

To present the matrix version, we need to clarify the definitions of inner product and Euclidean norm for matrices, both natural generalizations of those for vectors. For $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{m \times n}$,

$$\langle \mathbf{U}, \mathbf{V} \rangle = \sum_{i,j} u_{ij} v_{ij} \quad \text{and} \quad \|\mathbf{U}\|_F = \sqrt{\sum_{i,j} u_{ij}^2} = \sqrt{\langle \mathbf{U}, \mathbf{U} \rangle}. \quad (3.5)$$

In other words, let $\text{vec}(\mathbf{U})$ be the vectorized version of \mathbf{U} by sequentially stacking its columns into a long vector. We have

$$\langle \mathbf{U}, \mathbf{V} \rangle = \langle \text{vec}(\mathbf{U}), \text{vec}(\mathbf{V}) \rangle \quad \text{and} \quad \|\mathbf{U}\|_F = \|\text{vec}(\mathbf{U})\|_2. \quad (3.6)$$

Theorem 3.3 (Taylor's theorem: matrix-variable version). *Consider $f(\mathbf{X}) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$.*

– *If f is 1st-order differentiable at \mathbf{X}_0 , then*

$$f(\mathbf{X}_0 + \Delta) = \underbrace{f(\mathbf{X}_0) + \langle \nabla f(\mathbf{X}_0), \Delta \rangle}_{\text{first-order Taylor expansion}} + o(\|\Delta\|_F) \quad \text{as } \Delta \rightarrow \mathbf{0}. \quad (3.7)$$

– *If f is 2nd-order differentiable at \mathbf{X}_0 , then*

$$f(\mathbf{X}_0 + \Delta) = f(\mathbf{X}_0) + \underbrace{\langle \nabla f(\mathbf{X}_0), \Delta \rangle + \frac{1}{2} \langle \Delta, \nabla^2 f(\mathbf{X}_0)[\Delta] \rangle}_{\text{second-order Taylor expansion}} + o(\|\Delta\|_F^2) \quad \text{as } \Delta \rightarrow \mathbf{0}. \quad (3.8)$$

Here, $\nabla^2 f(\mathbf{X}_0)$ is a 4-dimensional tensor (i.e., 4-dimensional array) that collects all the second-order derivatives. It is possible to define tensor-matrix product, but it can quickly become overwhelming technical. For our purposes, we only need to know that: (1) $\nabla^2 f(\mathbf{X}_0)[\Delta]$ is an $m \times n$ matrix here; (2) this form is consistent with the vector-variable version, if we vectorize the variable: let $F(\text{vec}(\mathbf{X}_0)) \doteq f(\mathbf{X}_0)$, we have

$$\begin{aligned} F(\text{vec}(\mathbf{X}_0 + \Delta)) &= F(\text{vec}(\mathbf{X}_0)) + \langle \nabla F(\text{vec}(\mathbf{X}_0)), \text{vec}(\Delta) \rangle \\ &\quad + \frac{1}{2} \langle \text{vec}(\Delta), \nabla^2 F(\text{vec}(\mathbf{X}_0)) \text{vec}(\Delta) \rangle + o(\|\text{vec}(\Delta)\|_2^2). \end{aligned} \quad (3.9)$$

So, if we have to implement the second-order expansion for matrix-variable functions, a safe way is to vectorize the variable and then call the vector-variable version.

Now we want to put Taylor's theorems to good use. But before that, we need another important property of the Taylor expansion. In short, *Taylor expansion is unique*.

Theorem 3.4 (Asymptotic uniqueness of Taylor expansion—scalar version). *Let $f(x) : \mathbb{R} \rightarrow \mathbb{R}$ be k -times ($k \geq 1$ integer) differentiable at a point x_0 . If $P(\delta)$ is a k -th order polynomial satisfying $f(x_0 + \delta) - P(\delta) = o(\delta^k)$ as $\delta \rightarrow 0$, then $P(\delta) = f(x_0) + \sum_{i=1}^k \frac{1}{k!} f^{(k)}(x_0) \delta^k$, i.e., k -th order Taylor expansion.*

Why is this useful? Typically, we calculate derivatives and plug them into the Taylor formula to obtain the Taylor expansion. This theorem enables the reverse path. Suppose we have at hand a k -th order polynomial $P(\delta)$ satisfying $f(x_0 + \delta) = P(\delta) + o(|\delta|)$ —no matter how we obtain it, e.g., by the perturbation-expansion technique described above, we can compare it to the standard Taylor expansion form and read off the derivatives.

Example 3.5. Consider $f(x) = x^3$ and let us calculate $\nabla f(x)$ and $\nabla^2 f(x)$. For any infinitesimal perturbation δ ,

$$f(x + \delta) = (x + \delta)^3 = x^3 + 3x^2\delta + 3x\delta^2 + \delta^3. \quad (3.10)$$

First-order term in δ is $3x^2\delta$ and so $\nabla f(x) = 3x^2$. Second-order term in δ is $3x\delta^2$ and so $\nabla^2 f(x) = 6x$.

The uniqueness property also holds for the vector and matrix versions.

Theorem 3.6 (Asymptotic uniqueness of Taylor expansion—vector version). *Consider $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$.*

- *Assume $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ is 1st-order differentiable at \mathbf{x}_0 . If $P(\delta) \doteq f(\mathbf{x}_0) + \langle \mathbf{v}, \delta \rangle$ satisfies that $f(\mathbf{x}_0 + \delta) - P(\delta) \in o(\|\delta\|_2)$ as $\delta \rightarrow \mathbf{0}$, then $P(\delta) = f(\mathbf{x}_0) + \langle \nabla f(\mathbf{x}_0), \delta \rangle$, i.e., the 1st-order Taylor expansion, and $\mathbf{v} = \nabla f(\mathbf{x}_0)$;*

- Assume $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ is 2nd-order differentiable at \mathbf{x}_0 . If $P(\boldsymbol{\delta}) \doteq f(\mathbf{x}_0) + \langle \mathbf{v}, \boldsymbol{\delta} \rangle + \frac{1}{2} \langle \boldsymbol{\delta}, \mathbf{H} \boldsymbol{\delta} \rangle$ with \mathbf{H} symmetric satisfies that $f(\mathbf{x}_0 + \boldsymbol{\delta}) - P(\boldsymbol{\delta}) \in o(\|\boldsymbol{\delta}\|_2^2)$ as $\boldsymbol{\delta} \rightarrow \mathbf{0}$, then $P(\boldsymbol{\delta}) = f(\mathbf{x}_0) + \langle \nabla f(\mathbf{x}_0), \boldsymbol{\delta} \rangle + \frac{1}{2} \langle \boldsymbol{\delta}, \nabla^2 f(\mathbf{x}_0) \boldsymbol{\delta} \rangle$, i.e., the 2nd-order Taylor expansion, and $\mathbf{v} = \nabla f(\mathbf{x}_0)$, $\mathbf{H} = \nabla^2 f(\mathbf{x}_0)$.

The matrix version, as well as proofs of the asymptotic uniqueness properties and other forms of Taylor's theorems, can be found in Chapter 5 of [Col12]. Now we provide a couple of examples to show how the perturbation-expansion technique can help us derive gradient and Hessian from Taylor expansions.

Example 3.7. Let's consider $f(\mathbf{x}) = \|\mathbf{Ax} - \mathbf{b}\|_2^2$ again and try to derive $\nabla f(\mathbf{x})$ and $\nabla^2 f(\mathbf{x})$. From Example 2.4, we know that

$$f(\mathbf{x} + \boldsymbol{\delta}) = \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \underbrace{2 \langle \mathbf{Ax} - \mathbf{b}, \mathbf{A}\boldsymbol{\delta} \rangle}_{\text{1st-order in } \boldsymbol{\delta}} + \underbrace{\|\mathbf{A}\boldsymbol{\delta}\|_2^2}_{\text{2nd-order in } \boldsymbol{\delta}}. \quad (3.11)$$

To read off the gradient, we need to rearrange the 1-st order term into the form $\langle \clubsuit, \boldsymbol{\delta} \rangle$ for some \clubsuit . Now we need a useful rule for manipulating vector/matrix inner products.

Any leading matrix can be transposed and moved to the leading position of the other side of the inner product; similarly, any trailing matrix can be transposed and moved to the trailing position of the other side of the inner product. For example, consider matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ with compatible dimensions so that $\langle \mathbf{AB}, \mathbf{CD} \rangle$ is well defined. Then

$$\langle \mathbf{AB}, \mathbf{CD} \rangle = \langle \mathbf{B}, \mathbf{A}^\top \mathbf{CD} \rangle = \langle \mathbf{C}^\top \mathbf{AB}, \mathbf{D} \rangle = \langle \mathbf{A}, \mathbf{C} \mathbf{D} \mathbf{B}^\top \rangle = \langle \mathbf{ABD}^\top, \mathbf{C} \rangle. \quad (3.12)$$

This property can be derived from the cyclic property of matrix traces.

So we can rearrange the 1st order term as

$$2 \langle \mathbf{Ax} - \mathbf{b}, \mathbf{A}\boldsymbol{\delta} \rangle = \langle 2\mathbf{A}^\top(\mathbf{Ax} - \mathbf{b}), \boldsymbol{\delta} \rangle, \quad (3.13)$$

implying that $\nabla f(\mathbf{x}) = 2\mathbf{A}^\top(\mathbf{Ax} - \mathbf{b})$. For the 2nd order term,

$$\|\mathbf{A}\boldsymbol{\delta}\|_2^2 = \langle \mathbf{A}\boldsymbol{\delta}, \mathbf{A}\boldsymbol{\delta} \rangle = \langle \boldsymbol{\delta}, \mathbf{A}^\top \mathbf{A}\boldsymbol{\delta} \rangle, \quad (3.14)$$

which is to be compared to $\frac{1}{2} \langle \boldsymbol{\delta}, \nabla^2 f(\mathbf{x}) \boldsymbol{\delta} \rangle$, implying that $\nabla^2 f(\mathbf{x}) = 2\mathbf{A}^\top \mathbf{A}$.

Example 3.8. We now consider a matrix-variable problem with two blocks of variables

$$f(\mathbf{W}_1, \mathbf{W}_2) = \sum_i \|\mathbf{y}_i - \mathbf{W}_2 \mathbf{W}_1 \mathbf{x}_i\|_2^2 = \|\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}\|_F^2 \quad (3.15)$$

and try to derive the gradient. This is an objective corresponding to a two-layer linear neural network. Making infinitesimal perturbation to $\mathbf{W}_1, \mathbf{W}_2$, we obtain

$$f(\mathbf{W}_1 + \boldsymbol{\Delta}_1, \mathbf{W}_2 + \boldsymbol{\Delta}_2) = \|\mathbf{Y} - (\mathbf{W}_2 + \boldsymbol{\Delta}_2)(\mathbf{W}_1 + \boldsymbol{\Delta}_1) \mathbf{X}\|_F^2 \quad (3.16)$$

$$= \|(\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) - \mathbf{W}_2 \boldsymbol{\Delta}_1 \mathbf{X} - \boldsymbol{\Delta}_2 \mathbf{W}_1 \mathbf{X} - \boldsymbol{\Delta}_2 \boldsymbol{\Delta}_1 \mathbf{X}\|_F^2. \quad (3.17)$$

Now we need the fact: for matrices \mathbf{U}, \mathbf{V} of the same size, $\|\mathbf{U} + \mathbf{V}\|_F^2 = \|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2 + 2 \langle \mathbf{U}, \mathbf{V} \rangle$. So,

$$\begin{aligned} & \|(\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) - \mathbf{W}_2 \boldsymbol{\Delta}_1 \mathbf{X} - \boldsymbol{\Delta}_2 \mathbf{W}_1 \mathbf{X} - \boldsymbol{\Delta}_2 \boldsymbol{\Delta}_1 \mathbf{X}\|_F^2 = \\ & \quad \|(\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) - \mathbf{W}_2 \boldsymbol{\Delta}_1 \mathbf{X} - \boldsymbol{\Delta}_2 \mathbf{W}_1 \mathbf{X}\|_F^2 \\ & \quad + \underbrace{\|\boldsymbol{\Delta}_2 \boldsymbol{\Delta}_1 \mathbf{X}\|_F^2}_{o(\|\boldsymbol{\Delta}\|_F)} - \underbrace{2 \langle (\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) - \mathbf{W}_2 \boldsymbol{\Delta}_1 \mathbf{X} - \boldsymbol{\Delta}_2 \mathbf{W}_1 \mathbf{X}, \boldsymbol{\Delta}_2 \boldsymbol{\Delta}_1 \mathbf{X} \rangle}_{o(\|\boldsymbol{\Delta}\|_F)}. \end{aligned} \quad (3.18)$$

So we only need to focus on

$$\|(\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) - \mathbf{W}_2 \Delta_1 \mathbf{X} - \Delta_2 \mathbf{W}_1 \mathbf{X}\|_F^2 = \underbrace{\|\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}\|_F^2}_{f(\mathbf{W}_1, \mathbf{W}_2)} + \underbrace{\|\mathbf{W}_2 \Delta_1 \mathbf{X} + \Delta_2 \mathbf{W}_1 \mathbf{X}\|_F^2}_{o(\|\Delta\|_F)} - 2 \langle \mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}, \mathbf{W}_2 \Delta_1 \mathbf{X} + \Delta_2 \mathbf{W}_1 \mathbf{X} \rangle. \quad (3.19)$$

We now only need to compare the linear term $-2 \langle \mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}, \mathbf{W}_2 \Delta_1 \mathbf{X} + \Delta_2 \mathbf{W}_1 \mathbf{X} \rangle$ with

$$\left\langle \begin{bmatrix} \partial_{\mathbf{W}_1} f \\ \partial_{\mathbf{W}_2} f \end{bmatrix}(\mathbf{W}_1, \mathbf{W}_2), \begin{bmatrix} \Delta_1 \\ \Delta_2 \end{bmatrix} \right\rangle = \langle \nabla_{\mathbf{W}_1} f(\mathbf{W}_1, \mathbf{W}_2), \Delta_1 \rangle + \langle \nabla_{\mathbf{W}_2} f(\mathbf{W}_1, \mathbf{W}_2), \Delta_2 \rangle. \quad (3.20)$$

We have that

$$-2 \langle \mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}, \mathbf{W}_2 \Delta_1 \mathbf{X} \rangle = -2 \langle \mathbf{W}_2^\top (\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) \mathbf{X}^\top, \Delta_1 \rangle, \quad (3.21)$$

$$-2 \langle \mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}, \Delta_2 \mathbf{W}_1 \mathbf{X} \rangle = -2 \langle (\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) \mathbf{X}^\top \mathbf{W}_1^\top, \Delta_2 \rangle, \quad (3.22)$$

implying that

$$\nabla_{\mathbf{W}_1} f(\mathbf{W}_1, \mathbf{W}_2) = -2 \mathbf{W}_2^\top (\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) \mathbf{X}^\top, \quad (3.23)$$

$$\nabla_{\mathbf{W}_2} f(\mathbf{W}_1, \mathbf{W}_2) = -2 (\mathbf{Y} - \mathbf{W}_2 \mathbf{W}_1 \mathbf{X}) \mathbf{X}^\top \mathbf{W}_1^\top. \quad (3.24)$$

Note that in this example, due to Eq. (3.20), one can also perturb $\mathbf{W}_1(\mathbf{W}_1, \mathbf{W}_2)$ only (i.e., $\Delta_2 = \mathbf{0}$) to obtain $\nabla_{\mathbf{W}_1} f$ and similarly perturb \mathbf{W}_2 only to obtain $\nabla_{\mathbf{W}_2} f(\mathbf{W}_1, \mathbf{W}_2)$. This tends to make the process less messy. Similarly, for functions with multiple groups of variables, one can take turns to perturb one group each time to derive group-wise gradients (But, in general, this does NOT work for higher-order derivatives!).

$$f(x, y) = x^2 - y^2$$

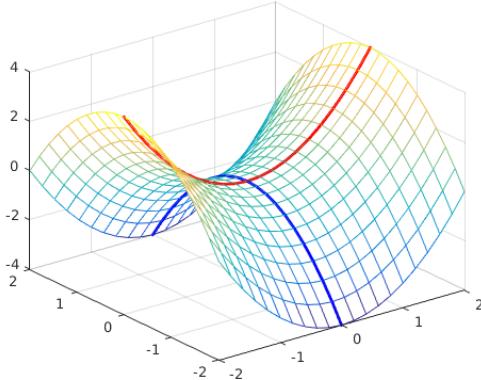


Figure 1: Blue: negative curvature; Red: positive curvature

Final words on this: we have discussed two or three techniques for deriving derivatives. For practical problems, often a mixture of these techniques works the best. So, stay flexible!

4 Directional derivatives and curvatures

Consider $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$.

- **directional derivative:** $D_{\mathbf{v}} f(\mathbf{x}_0) \doteq \frac{d}{dt} f(\mathbf{x}_0 + t\mathbf{v}) \Big|_{t=0}$, i.e., rate of change of f at \mathbf{x}_0 along the direction \mathbf{v}
- When f is 1-st order differentiable at \mathbf{x}_0 ,

$$D_{\mathbf{v}} f(\mathbf{x}_0) = \langle \nabla f(\mathbf{x}_0), \mathbf{v} \rangle.$$

• Now $D_{\mathbf{v}} f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ is another function. What is $D_{\mathbf{u}}(D_{\mathbf{v}} f)(\mathbf{x}_0)$? If f is 2nd-order differentiable at \mathbf{x}_0 ,

$$D_{\mathbf{u}}(D_{\mathbf{v}} f)(\mathbf{x}_0) = \langle \mathbf{u}, \nabla^2 f(\mathbf{x}_0) \mathbf{v} \rangle.$$

When $\mathbf{u} = \mathbf{v}$,

$$D_{\mathbf{u}}(D_{\mathbf{u}} f)(\mathbf{x}_0) = \langle \mathbf{u}, \nabla^2 f(\mathbf{x}_0) \mathbf{u} \rangle = \frac{d^2}{dt^2} f(\mathbf{x}_0 + t\mathbf{u}) \Big|_{t=0},$$

which is the **directional curvature** along u and grows quadratically with respect to $\|u\|_2$. To make it independent of the norm $\|u\|_2$, one can consider $\frac{\langle u, \nabla^2 f(x_0) u \rangle}{\|u\|_2^2}$.

The spectral property (i.e., the distribution of eigenvalues and eigenvectors) of $\nabla^2 f(x_0)$ determines the directional curvatures. In particular, eigenvector directions corresponding to negative (positive) eigenvalues of $\nabla^2 f(x_0)$ have negative (positive) curvatures.

Further reading

Chapters 3 & 5 of [DFO20] are particularly relevant, and you are encouraged to review the materials there.

Disclaimer

This set of notes is preliminary and has not been thoroughly proofread. Typos and factual errors are well expected, and hence use it with caution. Bug reports are very welcome and should be sent to Prof. Ju Sun via jusun@umn.edu.

References

- [Col12] Rodney Coleman, *Calculus on normed vector spaces*, Springer New York, 2012.
- [DFO20] Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, *Mathematics for machine learning*, Cambridge University Press, 2020.
- [Mun97] James R. Munkres, *Analysis on manifolds*, Taylor & Francis Inc, 1997.
- [Zor15] Vladimir A. Zorich, *Mathematical analysis I*, Springer Berlin Heidelberg, 2015.
- [Zor16] _____, *Mathematical analysis II*, Springer Berlin Heidelberg, 2016.