

Brief recap of matrix calculus

1. Matrix Calculus
 $\nabla f, \nabla^2 f$
2. Automatic Differentiation

Useful definitions and notations

We will treat all vectors as column vectors by default.

Matrix and vector multiplication

Let A be $m \times n$, and B be $n \times p$, and let the product AB be:



$$C = AB$$

$m \times p$ $m \times n$ $n \times p$

then C is a $m \times p$ matrix, with element (i, j) given by:

Наиболее простое решение $O(n^3)$

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$$



меньше, чем $O(n^2)$
не может быть

Алгоритмы
Шпрассена:
 $O(n^{\log_2 7})$


Let A be $m \times n$, and x be $n \times 1$, then the typical element of the product:

$$z = Ax$$

$m \times 1$ $m \times n$ $n \times 1$

$O(n^2)$

is given by:



$$z_k = a_k^T x$$

$$z_i = \sum_{k=1}^n a_{ik} x_k$$

Finally, just to remind:

$m \times p$ $m \times n$ $n \times p$ $p \times m$ $p \times n$ $n \times m$

• $C = AB$ $C^T = B^T A^T$

• $AB \neq BA$

• $e^A = \sum_{k=0}^{\infty} \frac{1}{k!} A^k$

Neural ODE

• $e^{A+B} \neq e^A e^B$ (but if A and B are commuting matrices, which means that $AB = BA$,
 $e^{A+B} = e^A e^B$)

• $\langle x, Ay \rangle = \langle A^T x, y \rangle$

$$\langle x, y \rangle = x^T y = y^T x = \langle y, x \rangle = \sum_{i=1}^n x_i \cdot y_i$$

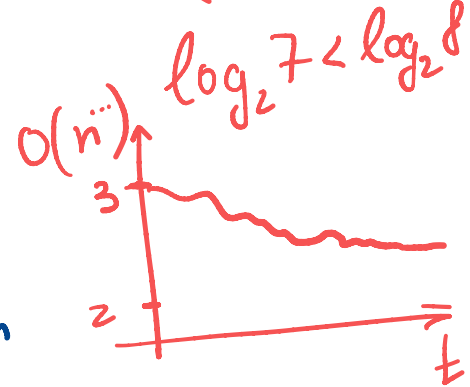
$$a \begin{pmatrix} 1 & 2 & 3 \\ 0 & -1 & 2 \end{pmatrix}$$

$$a^T b = 1 \cdot 0 + 2 \cdot (-1) + 3 \cdot 2 = 4$$

Gradient $\langle z^T, p^T B \rangle = \langle z^T B^T, p^T \rangle$

Let $f(x) : \mathbb{R}^n \rightarrow \mathbb{R}$, then vector, which contains all first order partial derivatives:

$$f'(x) = C^T x$$



named gradient of $f(x)$. This vector indicates the direction of the steepest ascent. Thus, vector $-\nabla f(x)$ means the direction of the steepest descent of the function in the point. Moreover, the gradient vector is always orthogonal to the contour line in the point.

General concept

The idea implies formulating a set of simple rules, which allows you to calculate derivatives just like in a scalar case. It might be convenient to use the differential notation here.

Differentials

$$df = f(x+dx) - f(x) \rightarrow 0$$

$x_k \in \mathbb{R}^n$
 $f(x_k - d\nabla f(x_k))$
 $d \downarrow \rightarrow 0$

$X, X+dx$
 $dx \rightarrow 0$
 $\|dx\| \rightarrow 0$


After obtaining the differential notation of df we can retrieve the gradient using the following formula:

$$df(x) = \langle \nabla f(x), dx \rangle$$

1. Градиент df
 2. Проверим к бугу $df = \langle \dots, dx \rangle \nabla f$

Then, if we have a differential of the above form and we need to calculate the second derivative of the matrix/vector function, we treat "old" dx as the constant dx_1 , then calculate $d(df)$

$$d^2 f(x) = \langle \nabla^2 f(x) dx_1, dx_2 \rangle = \langle H_f(x) dx_1, dx_2 \rangle$$



Properties

Let A and B be the constant matrices, while X and Y are the variables (or matrix functions).

- $dA = 0$
- $d(\alpha X) = \alpha(dX)$
- $d(AXB) = A(dX)B$
- $d(X + Y) = dX + dY$
- $d(X^T) = (dX)^T$
- $d(XY) = (dX)Y + X(dY)$
- $d\langle X, Y \rangle = \langle dX, Y \rangle + \langle X, dY \rangle$
- $d\left(\frac{X}{\phi}\right) = \frac{\phi dX - (d\phi)X}{\phi^2}$
- $d(\det X) = \det X \langle X^{-T}, dX \rangle$
- $d(\text{tr } X) = \langle I, dX \rangle$
- $df(g(x)) = \frac{df}{dg} \cdot dg(x)$
- $H = (J(\nabla f))^T$
- $d(X^{-1}) = -X^{-1}(dX)X^{-1}$

$A \quad B$
 $m \times n \quad m \times n$

$A^T B$
 $n \times m \quad m \times n$

$\langle A, B \rangle = \text{tr}(A^T B) = \text{tr}(B^T A)$

$\langle A, A \rangle = \|A\|_F^2$

$\|A\|_F = \sqrt{\sum_{ij} a_{ij}^2}$

$\langle A, A \rangle \neq \|A\|_F^2$

References

- Good introduction
- The Matrix Cookbook
- MSU seminars (Rus.)
- Online tool for analytic expression of a derivative.
- Determinant derivative

Examples

Example 1

$$df = \langle \dots, dx \rangle$$

\uparrow
 ∇f

Find $\nabla f(x)$, if $f(x) = \frac{1}{2}x^T Ax + b^T x + c$.

$$1. f = \frac{1}{2} \langle x, Ax \rangle + \langle b, x \rangle + c$$

$$2. df = d\left(\frac{1}{2} \langle x, Ax \rangle + \langle b, x \rangle + c\right) = d\left(\frac{1}{2} \langle x, Ax \rangle\right) + d(\langle b, x \rangle) + dc =$$

$$= \frac{1}{2} d(\langle x, Ax \rangle) + \langle db, x \rangle + \langle b, dx \rangle + dc =$$

$$= \frac{1}{2} (\langle dx, Ax \rangle + \langle x, d(Ax) \rangle) + \langle b, dx \rangle =$$

$$= \frac{1}{2} (\langle Ax, dx \rangle + \langle x, Adx \rangle) + \langle b, dx \rangle \equiv$$

$$\langle a, b \rangle = \langle b, a \rangle \quad \downarrow \langle Adx, x \rangle$$

$$\langle A^T x, dx \rangle$$

$$\nabla f = \frac{1}{2}(A + A^T)x + b$$

$$\equiv \langle \frac{1}{2}(A + A^T)x + b, dx \rangle \rightarrow$$

Example 2

Find $\nabla f(x)$, $f''(x)$, if $f(x) = -e^{-\langle x, x \rangle}$.

$$1. f = -e^{-\langle x, x \rangle}$$

$$2. df = d(-e^{-\langle x, x \rangle}) = -d(e^{-\langle x, x \rangle}) =$$

$$= -e^{-\langle x, x \rangle} \cdot d(-\langle x, x \rangle) = f \cdot (-1) (\langle dx, x \rangle + \langle x, dx \rangle) =$$

$$= -f \cdot \langle 2x, dx \rangle = \langle -2fx, dx \rangle$$

$$\Rightarrow \nabla f = 2e^{-\langle x, x \rangle} \cdot x$$

$$df = \langle \frac{1}{2}(A+A^T)x + b, dx_1 \rangle^{\text{const}}, \quad \nabla f = \frac{1}{2}(A+A^T)x + b$$

$$\nabla^2 f = ?$$

$$d(df) = \langle d(\frac{1}{2}(A+A^T)x + b), dx_1 \rangle =$$

$$= \langle \frac{1}{2}(A+A^T) dx_2, dx_1 \rangle =$$

$$= \langle \frac{1}{2}(A+A^T)^T dx_1, dx_2 \rangle \Rightarrow \boxed{\nabla^2 f = \frac{1}{2}(A+A^T)}$$

$$d^2 f = \langle \nabla^2 f dx_1, dx_2 \rangle$$

$$df = \langle -2fx, dx_1 \rangle$$

$$d^2 f = \langle \dots; dx_1, dx_2 \rangle$$

$$d(df) = d^2 f = \langle -2d(f \cdot x), dx_1 \rangle =$$

$$= \langle -2(df \cdot x + f \cdot dx), dx_1 \rangle = \langle -2(\langle -2fx, dx_1 \rangle x + f dx), dx_1 \rangle =$$

$$= \langle -2f (2\langle -x, dx_2 \rangle x + dx_2), dx_1 \rangle =$$

$$= -2f \langle -2x \cdot \langle x, dx_2 \rangle + dx_2, dx_1 \rangle =$$

$$= -2f \langle -2x x^T dx_2 + dx_2, dx_1 \rangle =$$

$$= -2f \langle (-2x x^T + I) dx_2, dx_1 \rangle =$$

$$= -2f \langle (I - 2x x^T)^T dx_1, dx_2 \rangle$$

$$(x x^T)^T = x^T \cdot x = x x^T$$

$n \times 1 \quad 1 \times n$

$$\Rightarrow \boxed{\nabla^2 f = -2f (I - 2x x^T)}$$

$$I dx_2 = dx_2$$

$$1 \cdot dx_2 = dx_2$$

$$f(x) = \text{tr}(AXB)$$

$$\nabla f = ?$$

$$df = \langle \dots, dX \rangle$$

trace cyclic property

$$\begin{aligned} \text{tr}(AXB) &= \\ &= \text{tr}(BA\underline{X}) = \\ &= \text{tr}(\underline{X}BA) \end{aligned}$$

$$\begin{aligned} f(x) &= \text{tr}(AXB) = \\ &= \text{tr}(BA\underline{X}) = \\ &= \langle (BA)^T, X \rangle \end{aligned}$$

$$\langle A, B \rangle = \text{tr}(A^T B)$$

$$df = \langle (BA)^T, dX \rangle$$

$$\nabla f = A^T B^T$$

$$f(x) = \text{tr}(X) \Leftrightarrow df =$$

$$\nabla f = ? \cdot \mathbf{1}_{I_n}$$

$$f: \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$$

$\nabla f \in \mathbb{R}^{n \times n} = \frac{\partial f}{\partial x_{ij}}$

$$\Leftrightarrow \text{tr}(I^T X) = \langle I_n, X \rangle$$

$$df = \langle I_n, dX \rangle \Rightarrow$$

$$\nabla f = I_n$$

$$f(x) = \frac{1}{2} \|Ax - b\|_2^2 =$$

$$= \frac{1}{2} \langle Ax - b, Ax - b \rangle$$

$df = ?$
 $\nabla f = ?$ $\langle x, x \rangle = \|x\|^2$
Linear Least Squares

$$\frac{1}{2} d(\langle Ax - b, Ax - b \rangle) =$$

$$\frac{1}{2} \cdot 2 \langle Ax - b, d(Ax - b) \rangle = \langle Ax - b, Adx \rangle =$$
$$= \langle A^T(Ax - b), dx \rangle$$

$$\nabla f = A^T(Ax - b)$$

$$f(x) = \ln(1 + \exp(\langle a, x \rangle)) = \sigma(\langle a, x \rangle)$$

$$df = \frac{d(1 + \exp(\langle a, x \rangle))}{1 + \exp(\langle a, x \rangle)} = \frac{\exp(\langle a, x \rangle) \cdot \langle a, dx \rangle}{1 + \exp(\langle a, x \rangle)}$$

$$\sigma(t) = \frac{1}{1 + e^{-t}} \quad \sigma'(t) = \sigma(t) \cdot (1 - \sigma(t))$$

$$\nabla f =$$

$$f(x) = \frac{1}{2} \|Ax - b\|^2$$

$$\nabla^2 f = ?$$

$$df = \langle A^T(Ax - b), dx_1 \rangle$$

$$d(df) = d(\langle A^T(Ax - b), dx_1 \rangle) = \langle d(A^T(Ax - b)), dx_1 \rangle =$$

$$= \langle A^T d(Ax - b), dx_1 \rangle = \langle A^T A dx, dx_1 \rangle = \langle dx_1, A^T A dx \rangle$$

$$\langle \nabla^2 f \cdot dx_1, dx \rangle = \langle A^T A dx_1, dx \rangle$$
$$(A^T A)^T = A^T A^T = A^T A$$
$$\nabla^2 f = A^T A$$

Example 3

$$\det X > 0$$

$$\ln x = \log_e x$$

Find $\nabla f(X)$, if $f(X) = \langle S, X \rangle - \log \det X$.

$$df = d(\langle S, X \rangle) - d(\ln \det X) = \langle S, dX \rangle - \frac{d(\det X)}{\det X}$$

$$= \langle S, dX \rangle - \frac{\cancel{\det X} \cdot \langle X^{-T}, dX \rangle}{\cancel{\det X}} = \langle S - X^{-T}, dX \rangle$$

$$\nabla f = S - X^{-T}$$

Example 4


Find $\nabla f(X)$, if $f(X) = \ln \langle Ax, x \rangle$, $A \in \mathbb{S}_{++}^n$

Automatic differentiation

$x_1 = W_1 \cdot x$ $x_2 = W_2 \cdot x_1$ NEURAL NETWORK


Idea

HEUTU BECA
 W_1, W_2, \dots, W_{228}

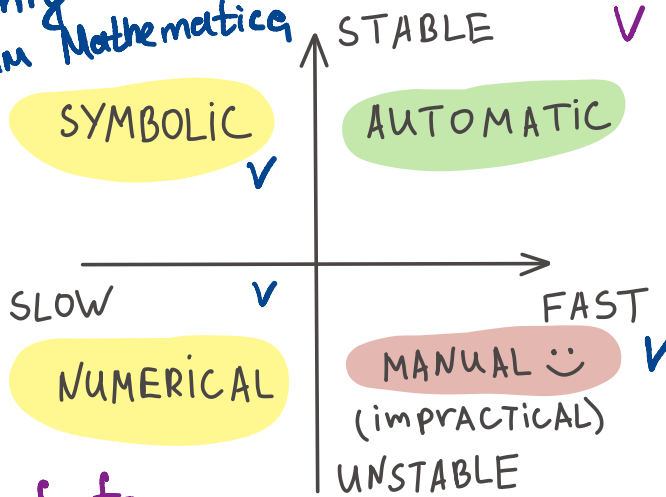


DIFFERENTIATION

SymPy
 Wolfram Mathematica



$f' = \frac{f_1 - f_0}{x_1 - x_0}$ $x_1 \rightarrow x_0$



$W_3^{k+1} = W_3^k - \alpha \frac{\partial L}{\partial W_3}$

$\frac{\partial L}{\partial W_1}, \frac{\partial L}{\partial W_2}, \dots$
 $\dots \frac{\partial L}{\partial W_{228}}$

Automatic differentiation is a scheme, that allows you to compute a value of the gradient of function with a cost of computing function itself only twice.

Chain rule

We will illustrate some important matrix calculus facts for specific cases

Univariate chain rule


Suppose, we have the following functions $R : \mathbb{R} \rightarrow \mathbb{R}$, $L : \mathbb{R} \rightarrow \mathbb{R}$ and $W \in \mathbb{R}$. Then

$$\frac{\partial R}{\partial W} = \frac{\partial R}{\partial L} \frac{\partial L}{\partial W} \quad W \rightarrow L \rightarrow R$$

Multivariate chain rule

The simplest example:

$$\frac{\partial}{\partial t} f(x_1(t), x_2(t)) = \frac{\partial f}{\partial x_1} \frac{\partial x_1}{\partial t} + \frac{\partial f}{\partial x_2} \frac{\partial x_2}{\partial t}$$



Now, we'll consider $f : \mathbb{R}^n \rightarrow \mathbb{R}$:

$$\frac{\partial}{\partial t} f(x_1(t), \dots, x_n(t)) = \frac{\partial f}{\partial x_1} \frac{\partial x_1}{\partial t} + \dots + \frac{\partial f}{\partial x_n} \frac{\partial x_n}{\partial t}$$

But if we will add another dimension $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, then the j -th output of f will be:

$$\frac{\partial}{\partial t} f_j(x_1(t), \dots, x_n(t)) = \sum_{i=1}^n \frac{\partial f_j}{\partial x_i} \frac{\partial x_i}{\partial t} = \sum_{i=1}^n J_{ji} \frac{\partial x_i}{\partial t},$$

СКАЛАРНАЯ
ЗАПИСЬ

where matrix $J \in \mathbb{R}^{m \times n}$ is the jacobian of the f . Hence, we could write it in a vector way:

$$\frac{\partial f}{\partial t} = J^T \frac{\partial x}{\partial t} \iff \left(\frac{\partial f}{\partial t} \right)^T = \left(\frac{\partial x}{\partial t} \right)^T J$$

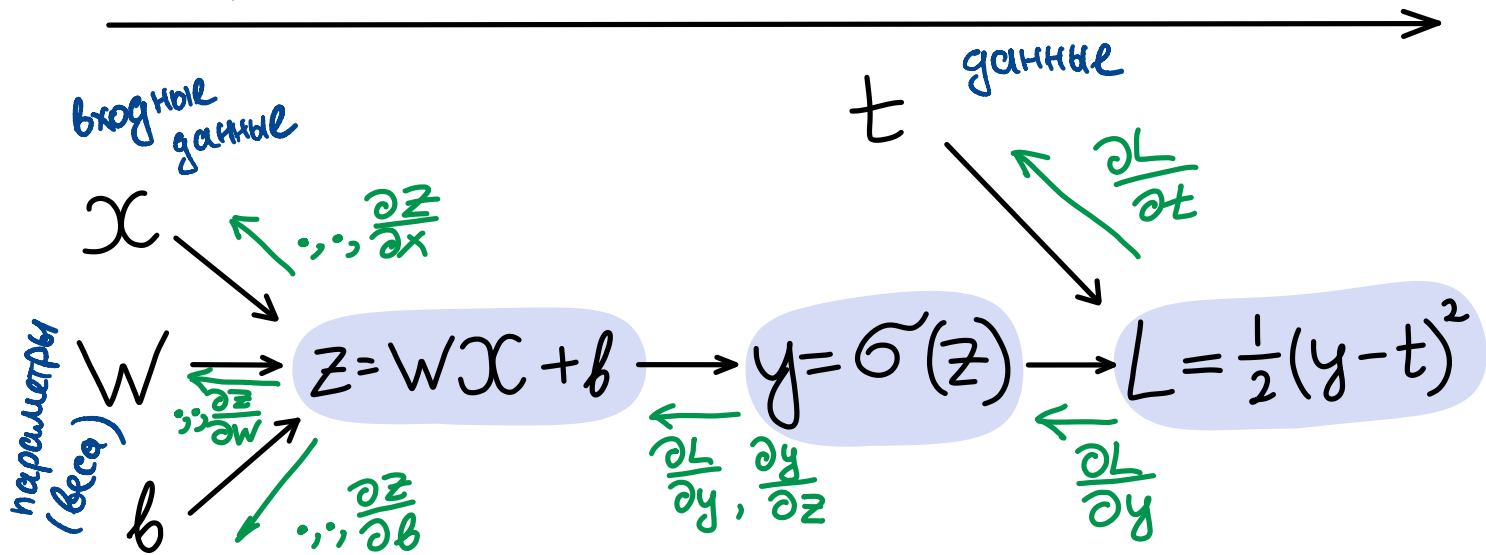
МАТРИЦНО-
ВЕКТОРНАЯ
ЗАПИСЬ

Backpropagation

The whole idea came from the applying chain rule to the computation graph of primitive operations

$$L = L(y(z(w, x, b)), t)$$

FORWARD PASS (COMPUTE LOSS)



BACKWARD PASS (compute derivatives)

Цель:

$$\frac{\partial L}{\partial w} = \frac{\partial L}{\partial y} \cdot \frac{\partial y}{\partial \sigma} \cdot \frac{\partial \sigma}{\partial z} \cdot \frac{\partial z}{\partial w}$$

$$z = wx + b \quad \frac{\partial z}{\partial w} = x, \quad \frac{\partial z}{\partial x} = w, \quad \frac{\partial z}{\partial b} = 0$$

$$y = \sigma(z) \quad \frac{\partial y}{\partial z} = \sigma'(z)$$

$$L = \frac{1}{2}(y - t)^2 \quad \frac{\partial L}{\partial y} = y - t, \quad \frac{\partial L}{\partial t} = t - y$$

All frameworks for automatic differentiation construct (implicitly or explicitly) computation graph. In deep learning, we typically want to compute the derivatives of the loss function L w.r.t. each intermediate parameter to tune them via gradient descent. For this purpose it is convenient to use the following notation:

узлы v_1, \dots, v_N

$$\overline{v_i} = \frac{\partial L}{\partial v_i}$$

Let v_1, \dots, v_N be a topological ordering of the computation graph (i.e. parents come before children). v_N denotes the variable we're trying to compute derivatives of (e.g. loss).

Forward pass:

- For $i = 1, \dots, N$:
 - Compute v_i as a function of its parents.

Backward pass:

- $\overline{v_N} = 1$
- For $i = N - 1, \dots, 1$:

- Compute derivatives $\bar{v}_i = \sum_{j \in \text{Children}(v_i)} \bar{v}_j \frac{\partial v_j}{\partial v_i}$

Note, that \bar{v}_j term is coming from the children of \bar{v}_i , while $\frac{\partial v_j}{\partial v_i}$ is already precomputed effectively.

~~$O(n^2)$~~ $O(n)$

Jacobian vector product

The reason why it works so fast in practice is that the Jacobian of the operations is already developed effectively in automatic differentiation frameworks. Typically, we even do not construct or store the full Jacobian, doing matvec directly instead.

BACKWARD PASS =
= 2x FORWARD PASS

Example: element-wise exponent

$$y = \exp(z) \quad J = \text{diag}(\exp(z)) \quad \bar{z} = \bar{y}J$$

See the examples of Vector-Jacobian Products from autodidact library:

```
defvjp(anp.add, lambda g, ans, x, y : unbroadcast(x, g),
        lambda g, ans, x, y : unbroadcast(y, g))
defvjp(anp.multiply, lambda g, ans, x, y : unbroadcast(x, y * g),
        lambda g, ans, x, y : unbroadcast(y, x * g))
defvjp(anp.subtract, lambda g, ans, x, y : unbroadcast(x, g),
        lambda g, ans, x, y : unbroadcast(y, -g))
defvjp(anp.divide, lambda g, ans, x, y : unbroadcast(x, g / y),
        lambda g, ans, x, y : unbroadcast(y, -g * x / y**2))
defvjp(anp.true_divide, lambda g, ans, x, y : unbroadcast(x, g / y),
        lambda g, ans, x, y : unbroadcast(y, -g * x / y**2))
```

Hessian vector product

Interesting, that a similar idea could be used to compute Hessian-vector products, which is essential for second order optimization or conjugate gradient methods. For a scalar-valued function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with continuous second derivatives (so that the Hessian matrix is symmetric), the Hessian at a point $x \in \mathbb{R}^n$ is written as $\partial^2 f(x)$. A Hessian-vector product function is then able to evaluate

$$v \mapsto \partial^2 f(x) \cdot v$$

for any vector $v \in \mathbb{R}^n$.

The trick is not to instantiate the full Hessian matrix: if n is large, perhaps in the millions or billions in the context of neural networks, then that might be impossible to store. Luckily, `grad` (in the `jax/autograd/pytorch/tensorflow`) already gives us a way to write an efficient Hessian-vector product function. We just have to use the identity

$$\partial^2 f(x)v = \partial[x \mapsto \partial f(x) \cdot v] = \partial g(x),$$

where $g(x) = \partial f(x) \cdot v$ is a new scalar-valued function that dots the gradient of f at x with the vector v . Notice that we're only ever differentiating scalar-valued functions of vector-valued arguments, which is exactly where we know `grad` is efficient.

```
import jax.numpy as jnp

def hvp(f, x, v):
    return grad(lambda x: jnp.vdot(grad(f)(x), v))(x)
```

Code

 [Open in Colab](#)

Materials

- [Autodidact](#) - a pedagogical implementation of Autograd
- [CSC321 Lecture 6](#)
- [CSC321 Lecture 10](#)
- [Why](#) you should understand backpropagation 😊
- [JAX autodiff cookbook](#)