

Как считать Hessien? $f: \mathbb{R}^n \rightarrow \mathbb{R}$

$$\nabla^2 f = \left(\frac{\partial^2 f}{\partial x_i \partial x_j} \right)_{ij}$$

\uparrow
 $\mathbb{R}^{n \times n}$

1.

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

Теперь $dx = dx_1 = \text{const}$

$$df = \langle \nabla f(x), dx_1 \rangle$$

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

Пример 1: $f(x) = \frac{1}{2} x^T A x - b^T x + c$

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

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

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

$$\left(\frac{1}{2}(A+A^T) \right)^T = \frac{1}{2}(A^T+A)$$

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

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

$n \times n$

$$H_f = f'' =$$

Пример 2 $f(x) = -e^{-x^T x}$ $\nabla f, \nabla^2 f = ?$

Решение:

$x^T x = \langle x, x \rangle$

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

1. $df = d(-e^{-\langle x, x \rangle}) = -\underbrace{e^{-\langle x, x \rangle}}_f \cdot d(-\langle x, x \rangle) =$

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

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

2. $d^2 f = d(\langle -2f x, dx \rangle) =$

$= \langle -2 d(fx), dx \rangle = -2 \langle \underbrace{df}_{n \times 1} \cdot x + f \underbrace{dx}_{n \times 1}, dx \rangle =$

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

$= -2 \langle -2f \cdot x^T dx \cdot x + f \cdot I dx, dx \rangle =$

$= -2f \langle -2 \cdot \underbrace{x}_{n \times 1} \cdot \underbrace{x^T}_{1 \times n} dx + \underbrace{I}_{n \times n} dx, dx \rangle =$

$= -2f \langle \underbrace{(-2x x^T + I)}_{n \times n} dx, dx \rangle$

$\underbrace{I - 2x x^T}_{n \times n}$

$\Rightarrow \nabla^2 f = -2f \cdot (I - 2x x^T)$

$\nabla^2 f = 2 e^{-\langle x, x \rangle} (I - 2x x^T)$

inner product

$x^T x$ - скаляр

$x x^T$ - матрица

outer product

$$\begin{pmatrix} x_1 x_1 & x_1 x_2 & \dots & x_1 x_n \\ \vdots & \vdots & \ddots & \vdots \\ x_n x_1 & \dots & \dots & x_n x_n \end{pmatrix}$$

$(I - 2x x^T)^T = I - 2(x x^T)^T =$

$= I - 2x x^T$

Задача:

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

$$x \in \mathbb{R}^n$$

$$A \in \mathbb{R}^{m \times n}$$

$$b \in \mathbb{R}^m$$

$$\nabla f = ?$$

$$\nabla^2 f = ?$$

квадрат
нормы
вектора

$$\|z\|_2^2 = \langle z, z \rangle$$

$$d\langle z, z \rangle = \langle 2z, dz \rangle$$

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

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

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

$$\underbrace{A^T A}_{n \times m \quad m \times n}_{n \times n}$$

$$\nabla^2 f \in \mathbb{R}^{n \times n}$$

$$\underbrace{AA^T}_{m \times m}$$

$$d(Ax - b) = d(Ax) - db = Adx - 0$$

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

$$d^2 f = \langle d(A^T Ax - A^T b), dx_i \rangle =$$

$$= \langle A^T A dx, dx_i \rangle =$$

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

$$(A^T A)^T = A^T A$$

$$\Rightarrow \nabla^2 f = A^T A$$

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

$x \in \mathbb{R}^n$
 $a \in \mathbb{R}^n$

$\nabla f, \nabla^2 f = ?$

$d(\ln f) = \frac{df}{f}$

$d(\exp(f)) = \exp(f) \cdot df$

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


$\Rightarrow \nabla f = \frac{\exp\langle a, x \rangle}{1 + \exp\langle a, x \rangle} \cdot a \Rightarrow$

$\frac{1}{1 + \exp(-\langle a, x \rangle)} = \sigma(\langle a, x \rangle)$

$\sigma'(x) = \sigma(x) \cdot (1 - \sigma(x))$

$\sigma(x) = \frac{1}{1 + e^{-x}}$

$d\sigma(f) = \sigma(f)(1 - \sigma(f)) \cdot df$



$$\nabla f = \underbrace{\sigma(\langle a, x \rangle)}_{\text{скаляр}} \cdot \underbrace{a}_{n \times 1}$$

$$d^2 f = \langle d(\sigma(\langle a, x \rangle) \cdot a), dx_i \rangle =$$

$$= \langle \sigma(\langle a, x \rangle) (1 - \sigma(\langle a, x \rangle)) \langle a, dx \rangle \cdot a, dx_i \rangle \quad d^2 f = \langle H dx, dx \rangle$$

$$= \sigma(1 - \sigma) \langle a^T dx a, dx_i \rangle =$$

$$H = H^T$$

$$= \sigma(1 - \sigma) \langle a a^T dx, dx_i \rangle =$$

$$H = \underbrace{\sigma(\langle a, x \rangle) (1 - \sigma(\langle a, x \rangle))}_{\text{скаляр}} \underbrace{a a^T}_{n \times 1 \cdot 1 \times n}$$

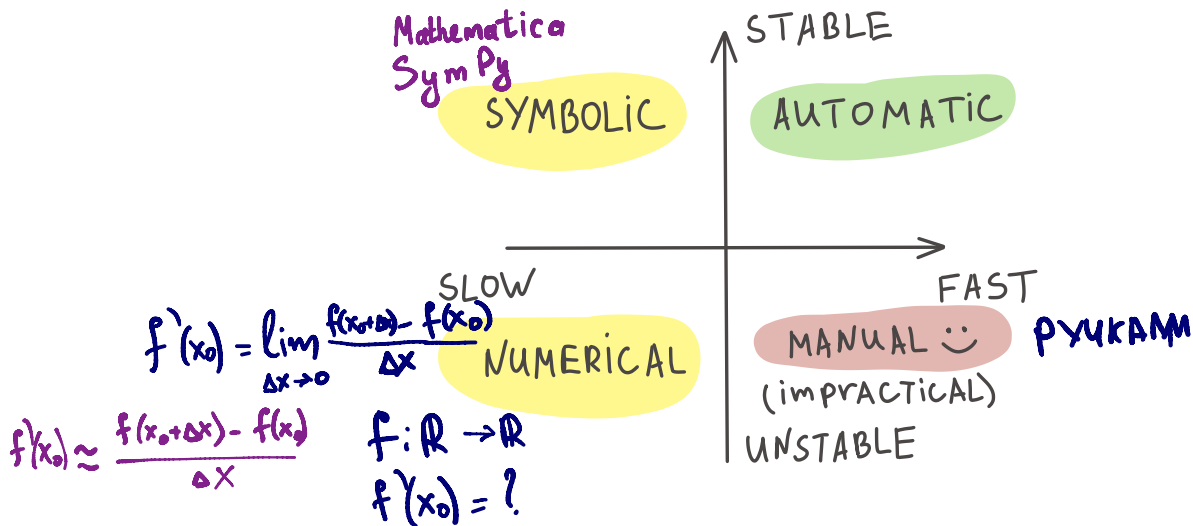
$$= \sigma(1 - \sigma) \langle a a^T dx, dx \rangle \Rightarrow$$

$$d^2 f = \langle \dots dx, dx \rangle$$

$$\nabla^2 f = \sigma(\langle a, x \rangle) (1 - \sigma(\langle a, x \rangle)) \cdot a a^T$$

Idea

DIFFERENTIATION



Automatic differentiation is a scheme, that allows you to compute a value of 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

$x \rightarrow W(x) \rightarrow L(W) \rightarrow R(L(W))$

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

произв. сложной функции


Multivariate chain rule

The simplest example:

$$\frac{\partial R}{\partial x} = \frac{\partial R}{\partial L} \frac{\partial L}{\partial w} \frac{\partial w}{\partial x}$$

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

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:

$$\begin{matrix} \vdots \\ \vdots \\ x_n(t) \end{matrix} \begin{matrix} \vdots \\ \vdots \\ f_m(x) \end{matrix} \quad \frac{\partial}{\partial t} f_j(x_1(t), \dots, x_n(t)) = \sum_{i=1}^n \frac{\partial f_j}{\partial x_i} \frac{\partial x_i}{\partial t} = \sum_{i=1}^n J_{ji} \frac{\partial x_i}{\partial t}, \quad J \in \mathbb{R}^{m \times n}$$

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

$$\begin{matrix} \frac{\partial f}{\partial t} \\ \frac{\partial f}{\partial t} \end{matrix} = \begin{matrix} J \\ J \end{matrix} \begin{matrix} \frac{\partial x}{\partial t} \\ \frac{\partial x}{\partial t} \end{matrix} \iff \left(\frac{\partial f}{\partial t} \right)^\top = \left(\frac{\partial x}{\partial t} \right)^\top J^\top$$

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

Backpropagation

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

рекуррентный список


$$W_{k+1} = W_k - \eta \frac{\partial L}{\partial W}(W_k)$$

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

подсчитывает функцию потерь

FORWARD PASS (COMPUTE LOSS)

прямой проход



BACKWARD PASS (compute derivatives)

обратный проход считаем производные

$$\begin{aligned} z &= wx + b & \frac{\partial z}{\partial w} &= x, \quad \frac{\partial z}{\partial x} = w, \quad \frac{\partial z}{\partial b} = 0 \\ y &= \sigma(z) & \frac{\partial y}{\partial z} &= \sigma'(z) \\ L &= \frac{1}{2}(y - t)^2 & \frac{\partial L}{\partial y} &= y - t, \quad \frac{\partial L}{\partial t} = t - y \end{aligned}$$

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 parameters in order to tune them via gradient descent. For this purpose it is convenient to use the following notation:

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

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).

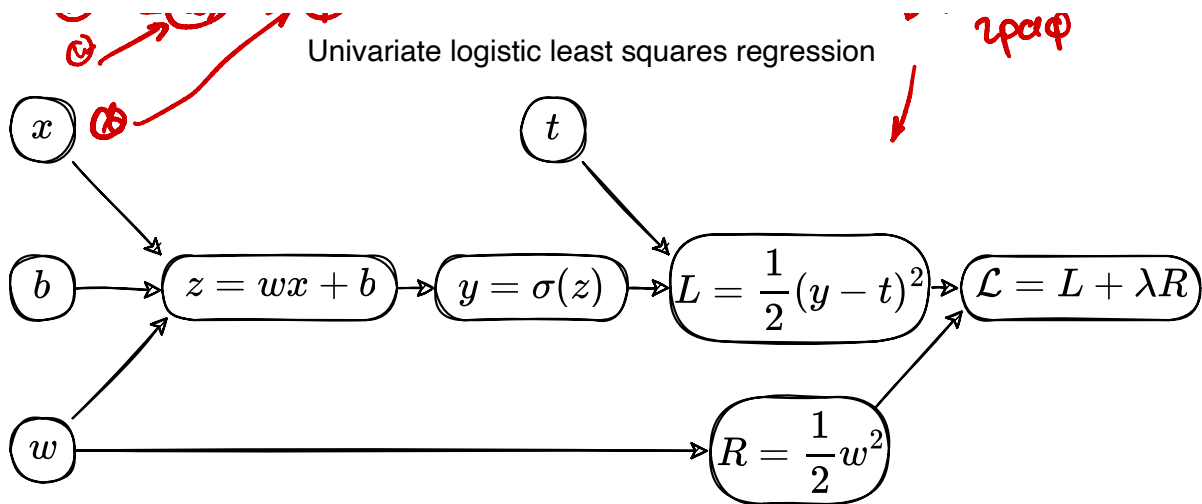
Forward pass:

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

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



Forward pass

$$z = wx + b$$

$$y = \sigma(z)$$

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

$$R = \frac{1}{2}w^2$$

$$\mathcal{L} = L + \lambda R$$

Backward pass

$$\bar{\mathcal{L}} = 1$$

$$\bar{R} = \bar{\mathcal{L}} \frac{d\mathcal{L}}{dR} = \bar{\mathcal{L}}\lambda$$

$$\bar{L} = \bar{\mathcal{L}} \frac{d\mathcal{L}}{dL} = \bar{\mathcal{L}}$$

$$\bar{y} = \bar{L} \frac{dL}{dy} = \bar{L}(y - t)$$

$$\bar{z} = \bar{y} \frac{dy}{dz} = \bar{y}\sigma'(z)$$

$$\bar{w} = \bar{z} \frac{dz}{dw} + \bar{R} \frac{dR}{dw} = \bar{z}x + \bar{R}w$$

$$\bar{b} = \bar{z} \frac{dz}{db} = \bar{z}$$

$$\bar{x} = \bar{z} \frac{dz}{dx} = \bar{z}w$$


Jacobian vector product

The reason why it works so fast in practice is that the Jacobian of the operations are already developed in effective manner in automatic differentiation frameworks.

Typically, we even do not construct or store the full Jacobian, doing matvec directly instead. Καυβιασ $O(n^2)$

Example: element-wise exponent

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



$y_i = e^{z_i}$
 $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$
 $\frac{\partial y_p}{\partial z_q}$
 $\frac{\partial y_1}{\partial z_1} = e^{z_1}$
 $\frac{\partial y_1}{\partial z_2} = 0$
 $\begin{pmatrix} e^{z_1} & 0 & 0 & 0 & 0 \\ 0 & e^{z_2} & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \ddots & e^{z_n} \end{pmatrix}$

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

```

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

$$\begin{pmatrix} e^{z_1} & 0 \\ \vdots & \vdots \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = \begin{pmatrix} e^{z_1} x_1 \\ \vdots \\ 0 \end{pmatrix}$$

```
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 the 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 vector-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)
```

1. Строим вычисл. граф примитивных арифм. операций. \oplus \otimes \div e^x

2. Каждой из операций сопоставляется J (произв. выхода по входу)
 якобианом

3. Во время backward pass собираются нужные производные.
4. Важно, что внутри реализовано

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](#)

Вычисление $f(x) \sim t_0$

Вычисление всех
произв. $\frac{\partial f}{\partial \dots}$ в выч. $\sim 2t$
графе