


Day 19 Differentiation by finite differences, cont'd. Differentiation arithmetics.

Let's begin by reviewing the error bound for the forward difference quotient in machine arithmetic, and comparing it with the empirically observed error ...



Using the methods we learned with Ch 1,
we obtained the leading order terms in the relative error are:

$$|\text{rel error in } f'| \approx \frac{1}{h} \left(1 + 2\mu \left| \frac{f}{f'} \right| \right) \frac{\epsilon_{mach}}{2} + \left(3 + \mu + \left| \frac{f''}{f'} \right| \right) \frac{\epsilon_{mach} h}{2} + \frac{h}{2} \left| \frac{f'''}{f'} \right|$$

← truncation error previously anticipated

Let us plot this bound along with the actual empirical error for the same example used to illustrate symbolic differentiation:

```

%matplotlib notebook
import sympy as sp
x = sp.symbols('x')
y = 3*x + sp.sin(7*x**2/(1+x)) # example from last time
yp = sp.diff(y,x).simplify()
ypp= sp.diff(y,x,x).simplify()
f = sp.lambdify(x,y,'numpy')
fp_exact = sp.lambdify(x,yp,'numpy')
fpp_exact = sp.lambdify(x,ypp,'numpy')

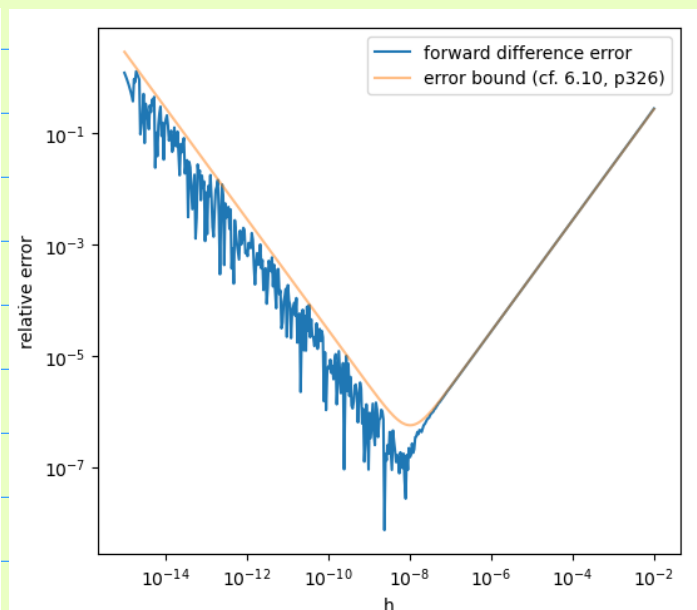
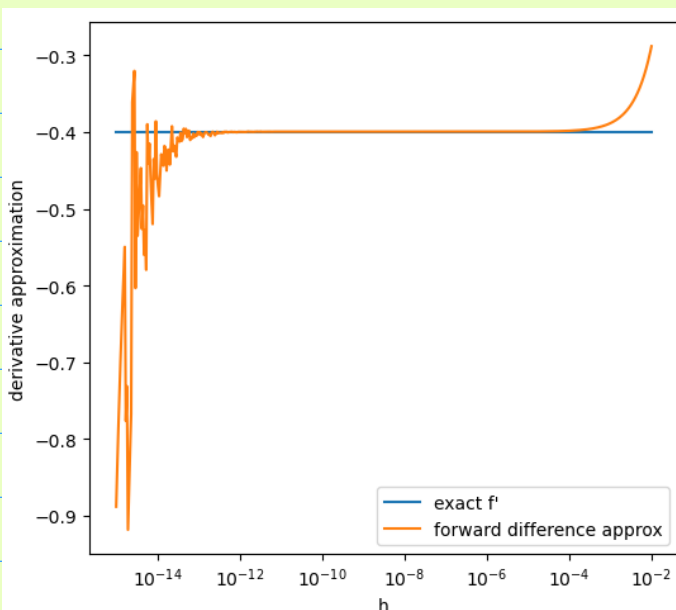
h = 10**np.linspace(-15,-2,500)
x = 1.1
fp_fd = (f(x+h)-f(x))/h # forward difference approximation to derivative of f for many different h values
error = np.abs(fp_fd - fp_exact(x))
relerr = error/np.abs(fp_exact(x))
import matplotlib.pyplot as plt
plt.figure(figsize=(6,12))
plt.subplot(211)
plt.semilogx(h,fp_exact(x)+0*h,label='exact f\''')
plt.semilogx(h,fp_fd,label='forward difference approx')
plt.legend()
plt.xlabel('h'); plt.ylabel('derivative approximation')
plt.subplot(212)
plt.loglog(h,relerr,label='forward difference error')
emach = 2**(-52)
M0 = np.abs(f(x))
M1 = np.abs(fp_exact(x))
M2 = np.abs(fpp_exact(x))

mu = 2 # estimate of multiple of emach/2 in error of evaluating f

relerr_bound = h/2*M2/M1 + emach/2*(1+2*mu*M0/M1)/h + (3+mu + M2/M1)*emach/2
plt.loglog(h,relerr_bound,label='error bound (cf. 6.10, p326)',alpha=0.5)

plt.legend()
plt.xlabel('h'); plt.ylabel('relative error');

```



* Zoom in to see how rounding errors vary with h in detail!

To find the optimal value of h , differentiate wrt h and set to 0:

$$\frac{d}{dh} \text{rel error bound} = -\frac{1}{h^2} \left(1 + 2\mu \left| \frac{f}{f'} \right| \right) \frac{\epsilon_{mach}}{2} + \frac{1}{2} \left| \frac{f''}{f'} \right| \stackrel{\text{set}}{=} 0$$

$$h = \sqrt{\frac{|f'|}{|f''|} \left(1 + 2\mu \left| \frac{f}{f'} \right| \right) \epsilon_{mach}} = \sqrt{\epsilon_{mach}} \sqrt{\frac{|f'|}{|f''|} \left(1 + 2\mu \left| \frac{f}{f'} \right| \right)}$$

$\frac{A}{h^2} = B \quad h = \sqrt{\frac{A}{B}}$

Summarizing, the error bound on the forward difference $\frac{f(x+h) - f(x)}{h} \approx f'(x)$ is minimized at

$$h_{opt} = \sqrt{\epsilon_{mach}}$$


some number depending on f , maybe 5

and from above

$$\text{rel. error} \Big|_{h=h_{opt}} = \text{Some number like 5} \sqrt{\epsilon_{mach}} \leftarrow \text{huge compared to } \epsilon_{mach}$$

because $\sqrt{\epsilon_{mach}} \gg \epsilon_{mach} !$

Can we do better? Yes - use a centered difference



If $f \in C^3$

$$f(x+h) = f(x) + hf'(x) + \frac{h^2}{2} f''(x) + \frac{h^3}{6} f'''(\xi_+)$$

$$f(x-h) = f(x) - hf'(x) + \frac{h^2}{2} f''(x) - \frac{h^3}{6} f'''(\xi_-)$$

Thus

$$\frac{f(x+h) - f(x-h)}{2h} = \frac{2hf'(x) + \frac{h^3}{3} (f'''(\xi_+) + f'''(\xi_-))}{2h}$$

$$= f'(x) + \frac{h^2}{6} f'''(\xi) = f'(x) + O(h^2)$$

Compare with the $O(h)$ error for the forward difference.

Let's do an empirical comparison of the errors in the forward and centered differences ...

Empirical observation of error in centered finite difference (and comparison with one-sided difference):

```
import numpy as np
def f(x): return np.cos(x)
def fp_exact(x): return -np.sin(x)
```


```
x = 1.1
```

```
h = 10**np.linspace(-15,-1,500)
```

```
fp_fd = (f(x+h)-f(x))/h #2 forward difference approximation to derivative of f for many different h values
fp_cd = (f(x+h)-f(x-h))/(2*h) # forward difference approximation to derivative of f for many different h values
```

```
error = np.abs(fp_fd - fp_exact(x))
error2 = np.abs(fp_cd - fp_exact(x))
```

```
import matplotlib.pyplot as plt
plt.loglog(h,error,label='forward difference')
plt.loglog(h,error2,label='centered difference')
plt.legend()
plt.xlabel('h'); plt.ylabel('error');
```



● "automatic" differentiation (AD), differentiation arithmetics

Would be great if we had another method that has

- (i) no truncation error
- (ii) no catastrophic round-off error

There is such a thing: we use a "differentiation arithmetic".

Suppose you have 2 functions f & g , and define $h = f \cdot g$ (their product).

Question: What do we need to know in order to compute say $h'(5.7)$?

Answer: $f(5.7), g(5.7), f'(5.7), g'(5.7)$

$$h'(5.7) = f'(5.7)g(5.7) + f(5.7)g'(5.7) \quad (\text{product rule})$$

So develop an arithmetic of objects like $\langle f(5.7), f'(5.7) \rangle$ $\langle g(5.7), g'(5.7) \rangle$,

for example $\langle 3.4, -1.3 \rangle, \langle 7, -2 \rangle$,

by defining the product of two objects like this:

$$\langle f_{\text{val}}, f_{\text{der}} \rangle * \langle g_{\text{val}}, g_{\text{der}} \rangle = \langle f_{\text{val}}g_{\text{val}}, f_{\text{der}}g_{\text{val}} + f_{\text{val}}g_{\text{der}} \rangle$$

Compare complex arithmetic:

$$\begin{aligned} & (f_{\text{re}} + i f_{\text{im}})(g_{\text{re}} + i g_{\text{im}}) \\ \langle f_{\text{re}}, f_{\text{im}} \rangle * \langle g_{\text{re}}, g_{\text{im}} \rangle & \\ = \langle f_{\text{re}}g_{\text{re}} - f_{\text{im}}g_{\text{im}}, f_{\text{re}}g_{\text{im}} + f_{\text{im}}g_{\text{re}} \rangle & \end{aligned}$$

Other arithmetic operations, and elementary transcendental functions, similarly ...

$$\underline{\pm} \quad \langle f_{\text{val}}, f_{\text{der}} \rangle \pm \langle g_{\text{val}}, g_{\text{der}} \rangle = \langle f_{\text{val}} \pm g_{\text{val}}, f_{\text{der}} \pm g_{\text{der}} \rangle$$

$$\langle f_{\text{der}}, f_{\text{val}} \rangle / \langle g_{\text{val}}, g_{\text{der}} \rangle = \langle f_{\text{val}}/g_{\text{val}}, \frac{g_{\text{val}}f_{\text{der}} - f_{\text{val}}g_{\text{der}}}{g_{\text{val}}^2} \rangle$$

$$\sin(\langle f_{\text{val}}, f_{\text{der}} \rangle) = \langle \sin(f_{\text{val}}), \cos(f_{\text{val}}) \cdot f_{\text{der}} \rangle$$

[space for other operations in a first-order single-variable differentiation arithmetic]

A series of horizontal blue lines spaced evenly down the page, providing a template for writing or calculations.

Example:

Let's do this example, first symbolically so we can check our AD answer.

An example for illustrating a Differentiation Arithmetic

```
x = sp.symbols('x')
y = x**2/(x-13)
y
```

$$\frac{x^2}{x-13}$$

```
yp = sp.diff(y,x)
yp
```

$$-\frac{x^2}{(x-13)^2} + \frac{2x}{x-13}$$

```
# evaluate y and y' at 3
y.subs({x:3}), yp.subs({x:3})    # use this to validate our AD answer
(-9/10, -69/100)
```

By AD: rules are:

$$\langle f, f' \rangle \pm \langle g, g' \rangle = \langle f \pm g, f' \pm g' \rangle$$
$$\langle f, f' \rangle * \langle g, g' \rangle = \langle fg, f'g + fg' \rangle$$
$$\langle f, f' \rangle / \langle g, g' \rangle = \langle f/g, (f'g - g'f)/g^2 \rangle$$

$x = \langle 3, 1 \rangle$

$$x^2 =$$

Now let's implement this Differentiation Arithmetic in code ...

Implementation of a differentiation arithmetic

with a gentle introduction to object-oriented programming in Python

```
import numpy as np

class ad:

    def __init__(self, val, der=0): # allows c = ad(3)
        self.val = val
        self.der = der

    def __repr__(self):
        return f'< {self.val}, {self.der} >'

    def __add__(self, other):
        return ad( self.val + other.val , self.der + other.der )

    def __sub__(self, other):
        return ad( self.val - other.val , self.der - other.der )

    def __mul__(self, other):
        return ad( self.val * other.val , self.der*other.val + self.val*other.der )

    def __truediv__(self, other):
        return ad( self.val / other.val , ( self.der*other.val - self.val*other.der )/ other.val**2 )

    def sin(self):
        return ad( np.sin( self.val ), np.cos(self.val)*self.der )

# to allow use of syntax sin(f)
def sin(x):
    if isinstance(x, ad):
        return x.sin()
    else:
        return np.sin(x)

# what we want to be able to do:

f = ad(3.4, -1.3)
g = ad(7, -2)
c = ad(3)
display(f)
display(g)
display(c)

f + g
print('product f*g')
display( f*g )
```


Check against symbolic differentiation

```
x = ad(3,1)
x*x
```

< 9, 6 >

```
x/x
```

< 1.0, 0.0 >

```
x + x
```

< 6, 2 >

```
x - x
```

< 0, 0 >

```
seven = ad(7,0)
seven*x
```

< 21, 7 >

```
z = ad(np.pi/2)
sin(z)
```

< 1.0, 0.0 >

```
x = sp.symbols('x')
y = 3*x + sp.sin(7*x**2/(1+x))
y
```

$$3x + \sin\left(\frac{7x^2}{x+1}\right)$$

```
yprime = sp.diff(y,x)
yprime
```

$$\left(-\frac{7x^2}{(x+1)^2} + \frac{14x}{x+1}\right) \cos\left(\frac{7x^2}{x+1}\right) + 3$$

```
y.subs({x:2.}), yprime.subs({x:2.})
```

(6.09131723555475, -3.19622486179360)

symbolic

```
x = ad(2.,1.)
three = ad(3,0)
seven = ad(7,0)
one = ad(1,0)
y = three*x + sin(seven*x*x/(one+x))y = 3*x + sp.sin(7*x**2/(1+x))
```

```
y
```

< 6.091317235554749, -3.1962248617936035 >

AD

agree

Extending the DA idea:

It's straightforward to extend the idea to

- (i) higher-order AD
- (ii) AD in multiple variables

(i) $\langle f, f', f'' \rangle * \langle g, g', g'' \rangle =$

$$\langle fg, f'g + fg', f''g + 2f'g' + fg'' \rangle$$

2nd der :: $(f'g + fg')' = f''g + f'g' + f'g' + fg''$ ↗

$x = \langle 3, 1, 0 \rangle$ $x^2 = x \cdot x = \langle 9, \underbrace{1 \cdot 3 + 3 \cdot 1}_6, \underbrace{0 \cdot 3 + 1 \cdot 1 + 1 \cdot 1 + 3 \cdot 0}_2 \rangle$

(ii) Suppose $f: \mathbb{R}^2 \rightarrow \mathbb{R}$

Represent f by $\langle f, \partial_1 f, \partial_2 f \rangle$

And the product rule would be $\uparrow f$

$$= \langle 9, 6, 2 \rangle$$

Symbolically, $y = x^2$ | $x=3$
 $y' = 2x$ | 6
 $y'' = 2$ | 2

$$\langle f, \partial_1 f, \partial_2 f \rangle * \langle g, \partial_1 g, \partial_2 g \rangle$$

$$= \langle fg, (\partial_1 f)g + f(\partial_1 g), (\partial_2 f)g + f(\partial_2 g) \rangle$$

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

represent f by $\langle f, \nabla f \rangle$

with product rule $\langle f, \nabla f \rangle * \langle g, \nabla g \rangle = \langle fg, \nabla f \cdot g + f \cdot \nabla g \rangle$

Machine learning, reverse-mode AD aka back-propagation

The above process is called "forward-mode" AD.
It is accurate to machine epsilon, and easy to program.
But it is not optimally efficient.

Machine learning - training a neural network - involves minimizing a real-valued function called the "loss", which is a measure of how poorly the network is performing the desired task.

The network is characterized by a large set of parameters called "weights",

$$w_0, w_1, w_2, \dots, w_{1000000000}$$

loss $y = N(\vec{w})$. Want to minimize y with respect to \vec{w} .
large vector of weights

The standard method is to move repeatedly in the direction of steepest descent. *gradient descent*

That direction is $-\nabla_{\vec{w}} N$

Because of the huge dimension of \vec{w} , we need to be as efficient as possible (as well as accurate).

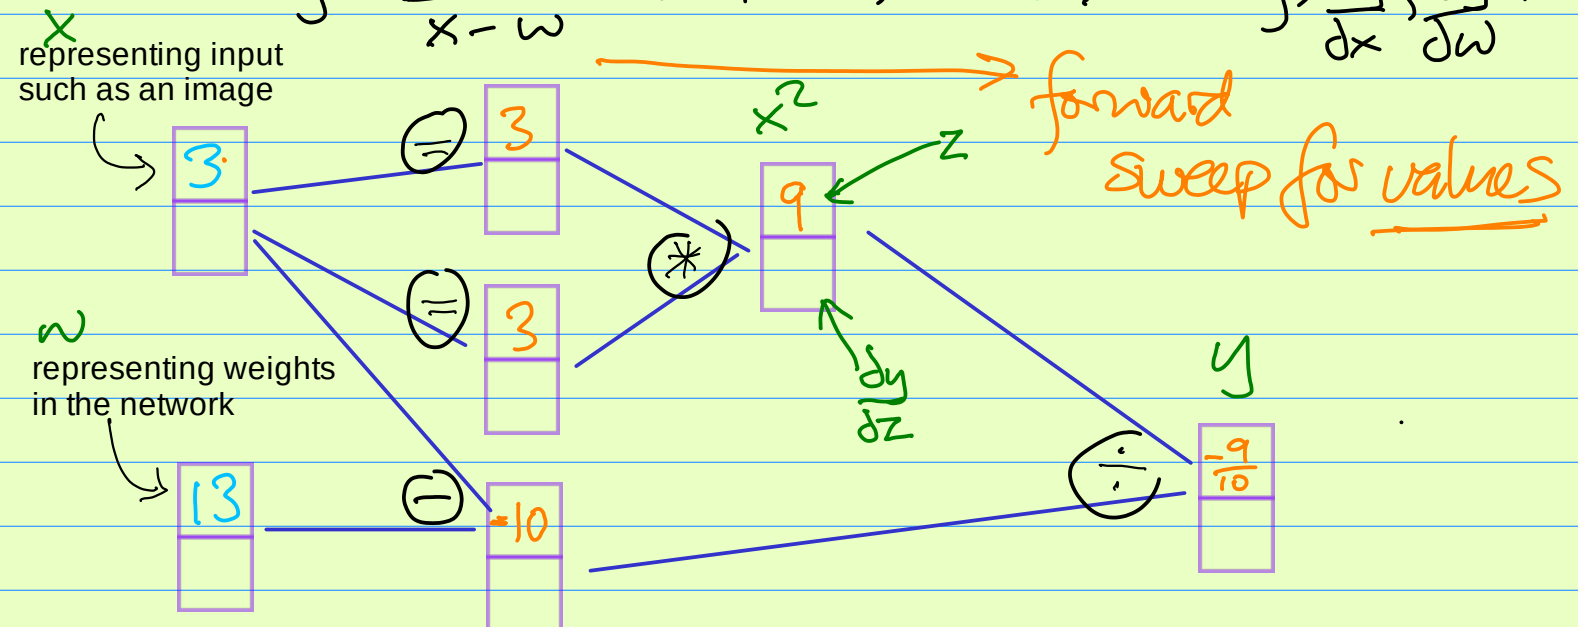
"Reverse-mode AD" or "back-propagation" is more efficient than forward-mode AD.

In the diagram below,
the upper slots are for the *values* of initial, intermediate, and final quantities in the calculation,
the lower slots are for the *derivative of the output value y with respect to the quantity*.

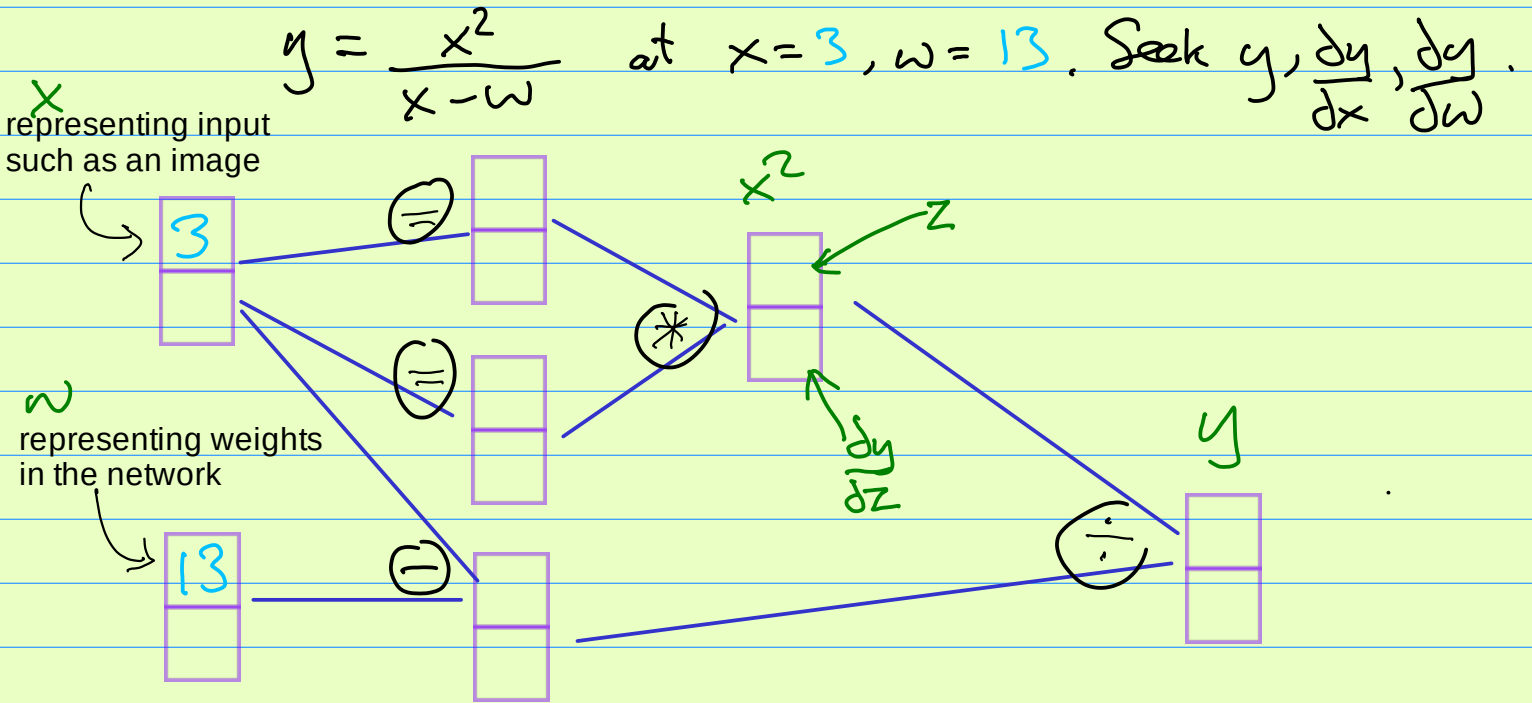
- Our goal is to fill in all the slots in two sweeps:
 (i) a forward sweep to fill in all the *values*
 (ii) a reverse sweep (back propagation) to fill in all the partial derivatives,

The last numbers we compute in the reverse sweep will be the desired components of the gradient of y .

Example: $y = \frac{x^2}{x-w}$ at $x=3, w=13$. Seek $y, \frac{dy}{dx}, \frac{dy}{dw}$.

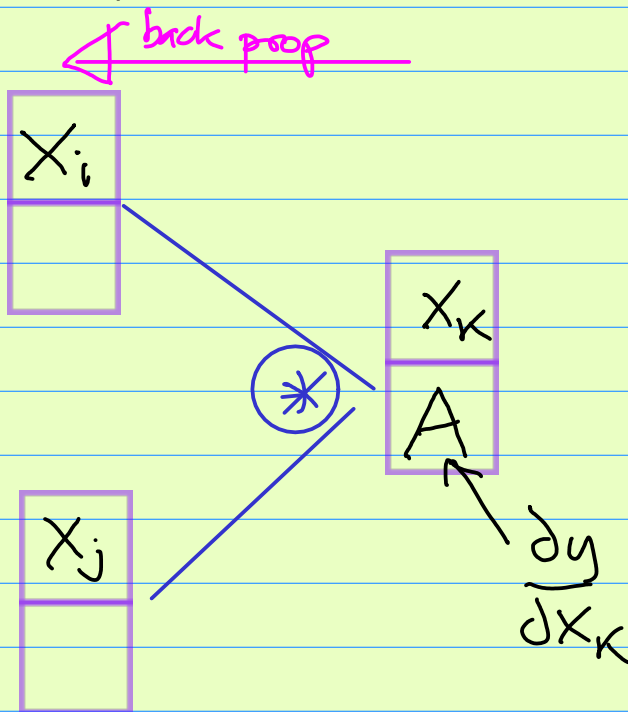


Forward sweep to fill in the values:



For the reverse sweep to fill in the partial derivatives of y , we need to develop some rules ...

Let's look at multiplication.



Chain rule is





$$\frac{dy}{dx_i} = \frac{dy}{dx_k} \frac{dx_k}{dx_i}$$

$=$

$=$

Exercise for you:

Find the back prop rules for the operations of \oplus , \ominus , \odot and branching.



Finally let's apply our back prop rules to compute the gradient of our example expression.

