

CSED 502: Computer Vision and Deep Learning

Tutorial: NumPy Fundamentals & Backpropagation

Welcome to class, we hope you've been enjoying the sun!

Reference Material

Rules of Broadcasting from Jake VanderPlas' *Python Data Science Handbook*:

- (1) If the two arrays differ in their number of dimensions, the shape of the one with fewer dimensions is padded with ones on its leading (left) side.
- (2) If the shape of the two arrays does not match in any dimension, the array with shape equal to 1 in that dimension is stretched to match the other shape.
- (3) If in any dimension the sizes disagree and neither is equal to 1, an error is raised.

Chain Rule for One Independent Variable:

Let $z = f(x, y)$ be a differentiable function. Further suppose that x and y are themselves differentiable functions of t , in other words $x = x(t)$ and $y = y(t)$. Then,

$$\frac{dz}{dt} = \frac{\partial z}{\partial x} \frac{dx}{dt} + \frac{\partial z}{\partial y} \frac{dy}{dt}$$

Chain Rule for Two Independent Variables:

Let $z = f(x, y)$ be a differentiable function, where x and y are themselves differentiable functions of a and b . In other words, $x = x(a, b)$ and $y = y(a, b)$. Then,

$$\frac{\partial z}{\partial a} = \frac{\partial z}{\partial x} \frac{\partial x}{\partial a} + \frac{\partial z}{\partial y} \frac{\partial y}{\partial a}$$

and

$$\frac{\partial z}{\partial b} = \frac{\partial z}{\partial x} \frac{\partial x}{\partial b} + \frac{\partial z}{\partial y} \frac{\partial y}{\partial b}$$

Generalized Chain Rule:

Let $w = f(x_1, x_2, \dots, x_m)$ be a differentiable function of m independent variables, and let $x_i = x_i(t_1, t_2, \dots, t_n)$ be a differentiable function of n independent variables. Then,

$$\frac{\partial w}{\partial t_j} = \frac{\partial w}{\partial x_1} \frac{\partial x_1}{\partial t_j} + \frac{\partial w}{\partial x_2} \frac{\partial x_2}{\partial t_j} + \dots + \frac{\partial w}{\partial x_m} \frac{\partial x_m}{\partial t_j}$$

for any $j \in 1, 2, \dots, n$.

Intuition for Backprop

Recall some basic facts:

- 1) The loss function L measures how "bad" our current model is.
- 2) L is a function of our parameters W .
- 3) We want to minimize L .

Thus, we update W to minimize L using $\frac{\partial L}{\partial W}$.

For example, if $\frac{\partial L}{\partial W_1}$ was positive, increasing W_1 would increase L . Accordingly, we'd choose to decrease W_1 .

More generally, `weights += (-1 * step_size * gradient)`.

Unfortunately, taking the derivative $\frac{\partial L}{\partial W}$ can get extremely difficult, especially at the scale of state-of-the-art models. For instance, GLM-4.5 has 92 hidden layers and 32 billion parameters. Imagine taking 32 billion derivatives, with each derivative having hundreds of applications of chain rule.

Instead, we employ a technique known as **backprop**.

First, we split our function into multiple equations until there is *one operation per equation*. This process is known as **staged computation**. Next, we take the derivatives of each of these smaller equations, before finally linking them together using **chain rule**.

Common Gates

Feel free to take notes on the common backprop gates here.

1. Dimension: Impossible

Determine if NumPy allows the **addition** of the following pairs of arrays, and if applicable determine what the result's dimensions will be.

(a) Where `x.shape` is $(2,)$ and `y.shape` is $(2, 1)$

(b) Where `x.shape` is $(4,)$ and `y.shape` is $(4, 1, 1)$

(c) Where `x.shape` is $(4, 2)$ and `y.shape` is $(2, 4, 1)$

(d) Where `x.shape` is $(8, 3)$ and `y.shape` is $(2, 8, 1)$

(e) Where `x.shape` is $(6, 5, 3)$ and `y.shape` is $(6, 5)$

Determine if NumPy allows the **matrix multiplication** of the following pairs of arrays, and if applicable determine what the result's dimensions will be.

(f) Where `a.shape` is (5, 4) and `b.shape` is (4, 8).

(g) Where `a.shape` is (3, 5, 4) and `b.shape` is (3, 4, 8).

(h) Where `a.shape` is (3, 5, 4) and `b.shape` is (5, 4, 8).

(i) Where `a.shape` is (1, 5, 4) and `b.shape` is (5, 4, 8).

(j) Where `a.shape` is (2, 5, 4) and `b.shape` is (3, 2, 4, 8).

2. The More (Derivatives) The Merrier

(a) Let $z = 2x + y$, with $x = \ln(t)$ and $y = \frac{1}{3}t^3$. Find $\frac{dz}{dt}$.

(b) Let $z = x^2y - y^2$ where $x = t^2$ and $y = 2t$. Find $\frac{dz}{dt}$. Your answer should be in terms of t .

(c) Let $z = 3x^2 - 2xy + y^2$. Also let $x = 3a + 2b$ and $y = 4a - b$. Find $\frac{\partial z}{\partial a}$ and $\frac{\partial z}{\partial b}$.

(d) Let $w = f(x, y, z)$, $x = x(t, u, v)$, $y = y(t, u, v)$ and $z = z(t, u, v)$. Find the formula for $\frac{\partial w}{\partial t}$.

3. Compute and Conquer

For each function below, use the staged computation approach to split it into smaller equations.

(a) $f(x, y, z) = (x + y)z$

(b) $h(x, y, z) = (x^2 + 2y)z^3$

(c) $g(x, y, z) = (\ln(x) + \sin(y))^2 + 4x$

4. Oh, node way!

For each function below:

- (i) construct a computational graph
- (ii) do a forward and backward pass through the graph using the provided input values
- (iii) complete the Python function for a combined forward and backward pass

It may be useful to consider how you split these functions into smaller equations in the question above.

(a) $f(x, y, z) = (x + y)z$ with input values $x = 1, y = 3, z = 2$

```
1  import numpy as np
2
3  # inputs: NumPy arrays `x`, `y`, `z` of identical size
4  # outputs: forward pass in `out`, gradients for x, y, z in `fx`, `fy`, `fz` respectively
5  def q2a(x, y, z):
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20      return out, fx, fy, fz
```

Ignore the line numbers, they do NOT correspond to the number of lines you need to write.

(b) $h(x, y, z) = (x^2 + 2y)z^3$ with input values $x = 3, y = 1, z = 2$

```
1  import numpy as np
2
3  # inputs: NumPy arrays `x`, `y`, `z` of identical size
4  # outputs: forward pass in `out`, gradients for x, y, z in `hx`, `hy`, `hz` respectively
5  def q2b(x, y, z):
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28      return out, hx, hy, hz
```

Ignore the line numbers, they do NOT correspond to the number of lines you need to write.

(c) $g(x, y, z) = (\ln(x) + \sin(y))^2 + 4x$ with input values $x = e, y = \frac{\pi}{2}, z = 2$

Python function printed on the following page.

```
1 import numpy as np
2
3 # inputs: NumPy arrays `x`, `y`, `z` of identical size
4 # outputs: forward pass in `out`, gradients for x, y, z in `gx`, `gy`, `gz` respectively
5 def q2c(x, y, z):
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50     return out, gx, gy, gz
```

Ignore the line numbers, they do NOT correspond to the number of lines you need to write.

5. Sigmoid Shenanigans

Consider the Sigmoid activation function:

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

Draw a computational graph and work through the backpropagation. Then, fill in the Python function. If you finish early, work through the analytical derivative for Sigmoid.

As a hint, you could split Sigmoid into the following functions:

$$a(x) = -x \quad b(x) = e^x \quad c(x) = 1 + x \quad d(x) = \frac{1}{x}$$

Observe that chaining these operations gives us Sigmoid: $d(c(b(a(x)))) = \sigma(x)$.

Suppose $x = 2$. What would the gradient with respect to x be? Feel free to use a calculator on this part.

You should have gotten around 0.1. If the step size is 0.2, what would the value of x be after taking one gradient descent step? As a hint, remember that `parameters -= step_size * gradient`.

```
1 import numpy as np
2
3 # inputs:
4 #   - a NumPy array `x`
5 # outputs:
6 #   - `out`: the result of the forward pass
7 #   - `fx` : the result of the backwards pass
8 def sigmoid(x):
9     # provided: forward pass with cache
10    a = -x
11    b = np.exp(a)
12    c = 1 + b
13    d = c ** -1
14    out = d
15
16    # TODO: backwards pass, "fx" represents  $df / dx$ 
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40    return out, fx
```

Ignore the line numbers, they do NOT correspond to the number of lines you need to write.

6. A Backprop a Day Keeps the Derivative Away

Consider the following function:

$$f = \frac{\ln x \cdot \sigma(\sqrt{y})}{\sigma((x+y)^2)}$$

Break the function up into smaller parts, then draw a computational graph and finish the Python function.

For reference, the derivative of Sigmoid is $\sigma(x) \cdot (1 - \sigma(x))$.

The TA solution breaks the function into 8 additional equations and rewrites f in terms of 2 of those additional equations. Yours doesn't have to match this exactly.

Python function printed on the following page.

```

1  import numpy as np
2
3  # helper function
4  def sigmoid(x):
5      return 1/(1 + np.exp(-x))
6
7  # inputs: NumPy arrays `x`, `y`
8  # outputs: forward pass in `out`, gradient for x in `fx`, gradient for y in `fy`
9  def complex_layer(x, y):
10
11      # forward pass
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29      # backwards pass
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58      return out, fx, fy

```

Ignore the line numbers, they do NOT correspond to the number of lines you need to write.

7. Vector Virtuosity

Consider the following function,

$$f(W, x) = \|W \cdot x\|^2 = \sum_{i=1}^n (W * x)_i^2$$

where $W \in \mathbb{R}^{n \times n}$ and $x \in \mathbb{R}^n$.

First draw the function's computation graph. Then compute the forward pass for the following inputs.

$$W = \begin{bmatrix} 0.1 & 0.5 \\ -0.3 & 0.8 \end{bmatrix} \quad x = \begin{bmatrix} 0.2 \\ 0.4 \end{bmatrix}$$

Lastly, compute the backward pass. Verify your answer by deriving the closed forms of $\nabla_W f$ and $\nabla_x f$.

This page intentionally left blank. Please use it for notes or scratch work.

This page intentionally left blank. Please use it for notes or scratch work.