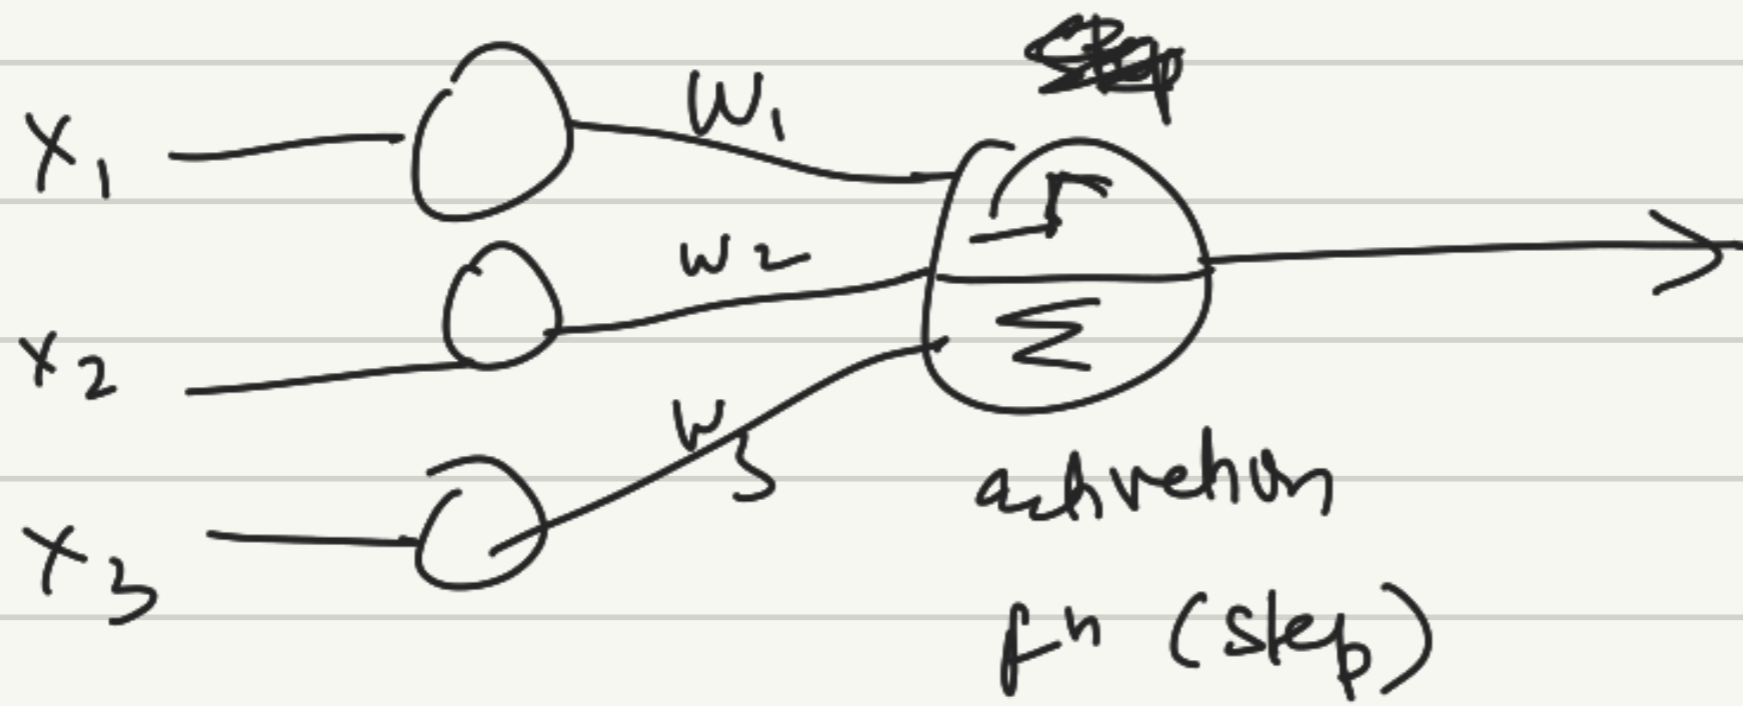


# Dec 29 - 2025

ANN - Artificial neural net

TLU



output:  
 $h_{w,b}(x) = \text{step}(W^T x + b)$

$$X = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \quad W = \begin{pmatrix} w_1 \\ w_2 \\ w_3 \end{pmatrix}$$

Training a neuron

### Hebb's rule

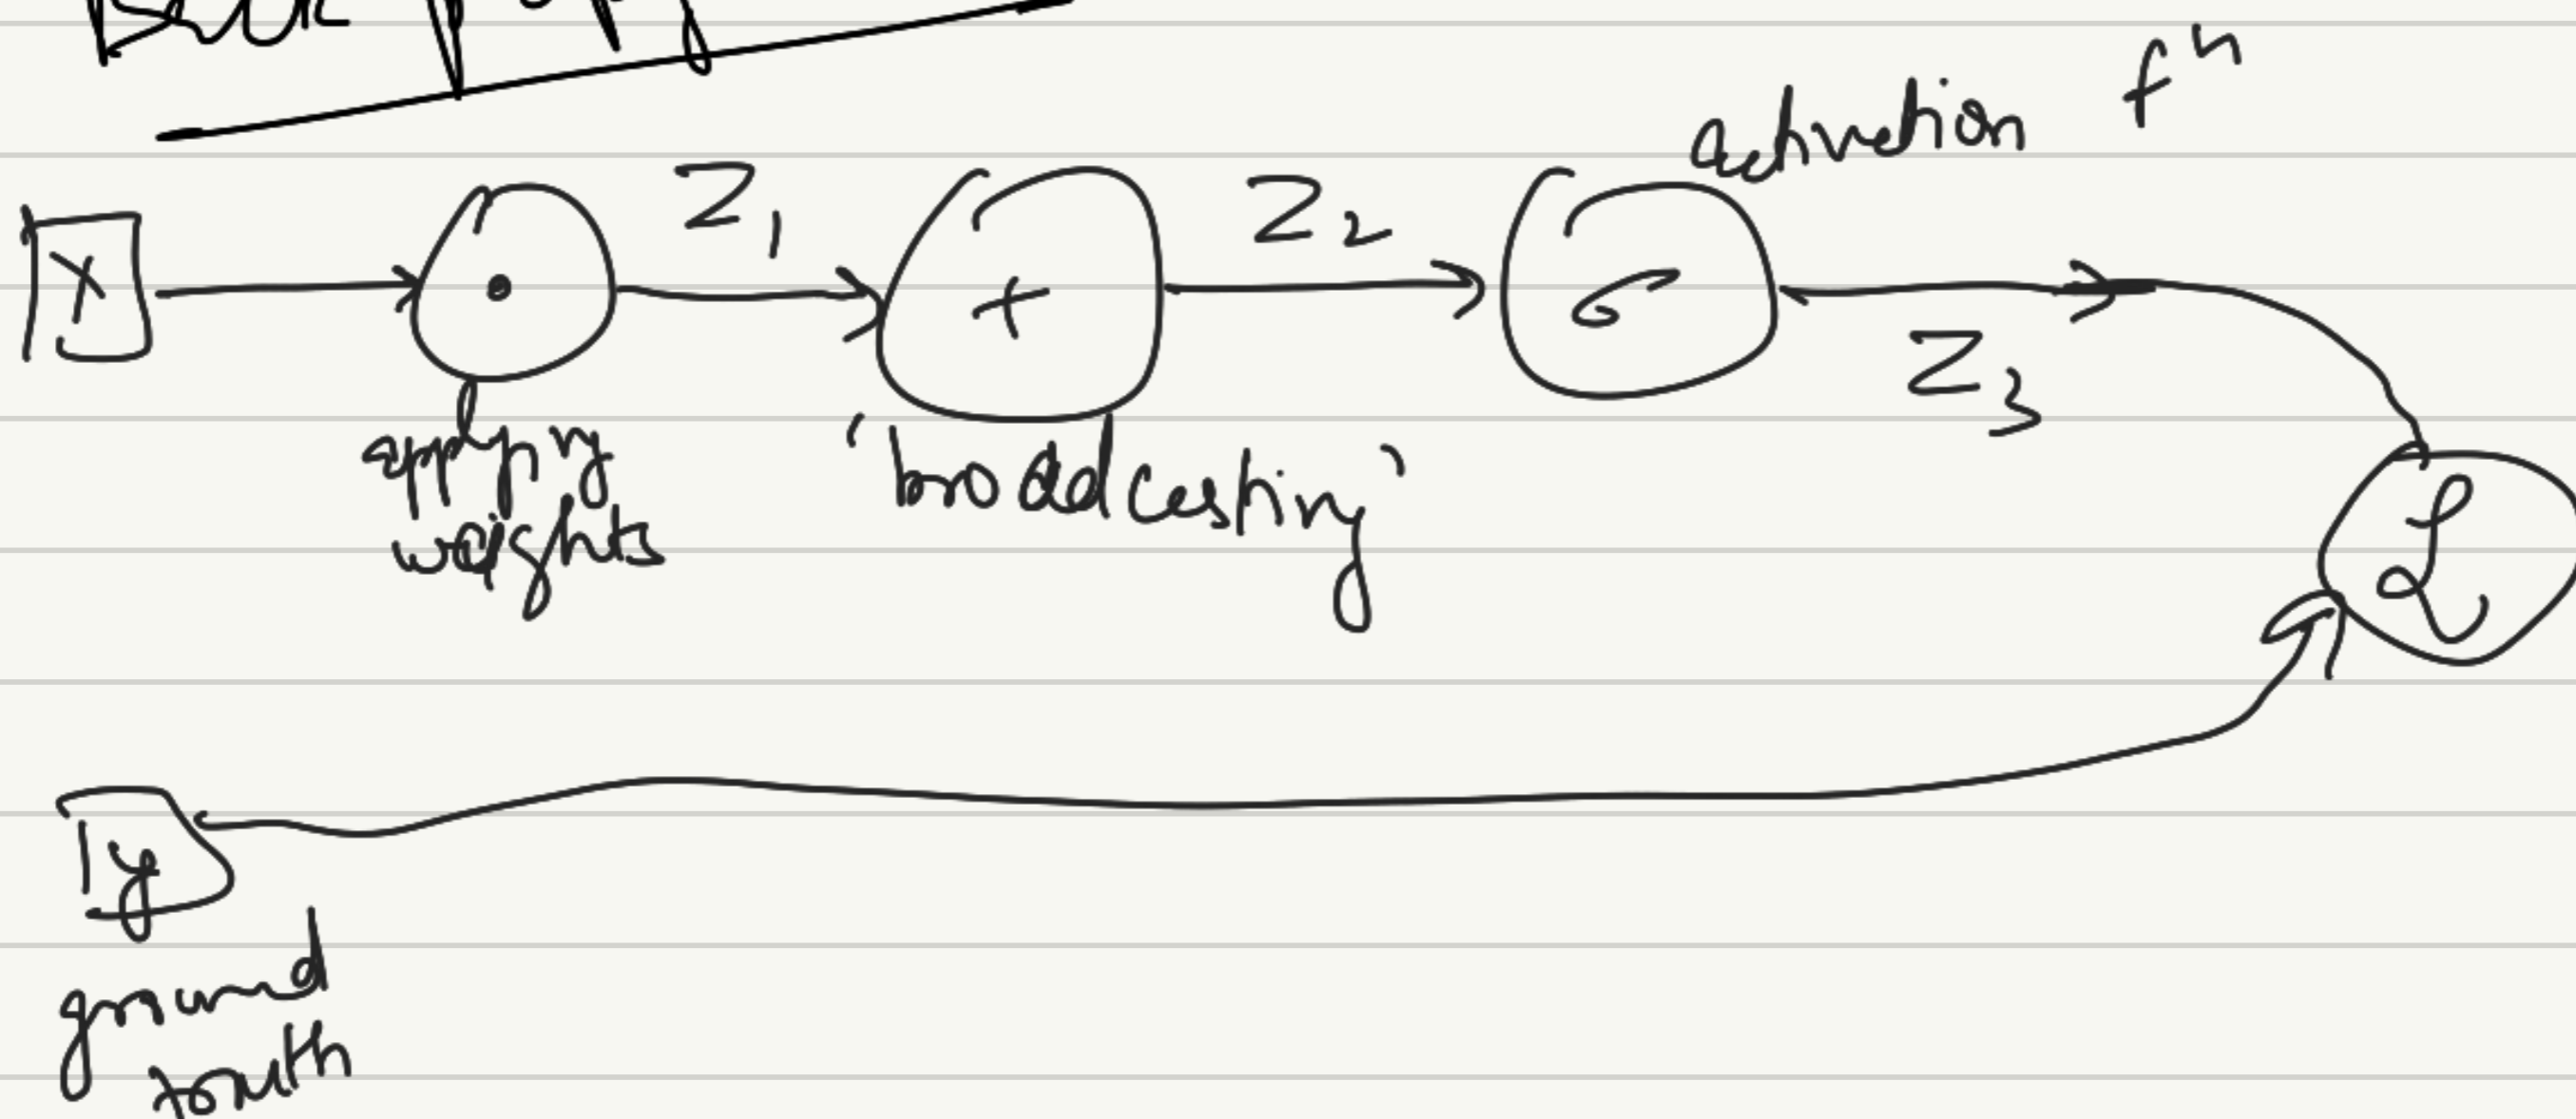
$$w_{ij} \text{ (next step)} = w_{ij} + \eta (y_i - \hat{y}_i) x_i$$

learning rate

loss fn:  $L = (y - z_3)^2$

$\uparrow$  ground truth       $\uparrow$  ML predicted

### Back propagation



calculate loss:  
 $L = (y - z_3)^2$

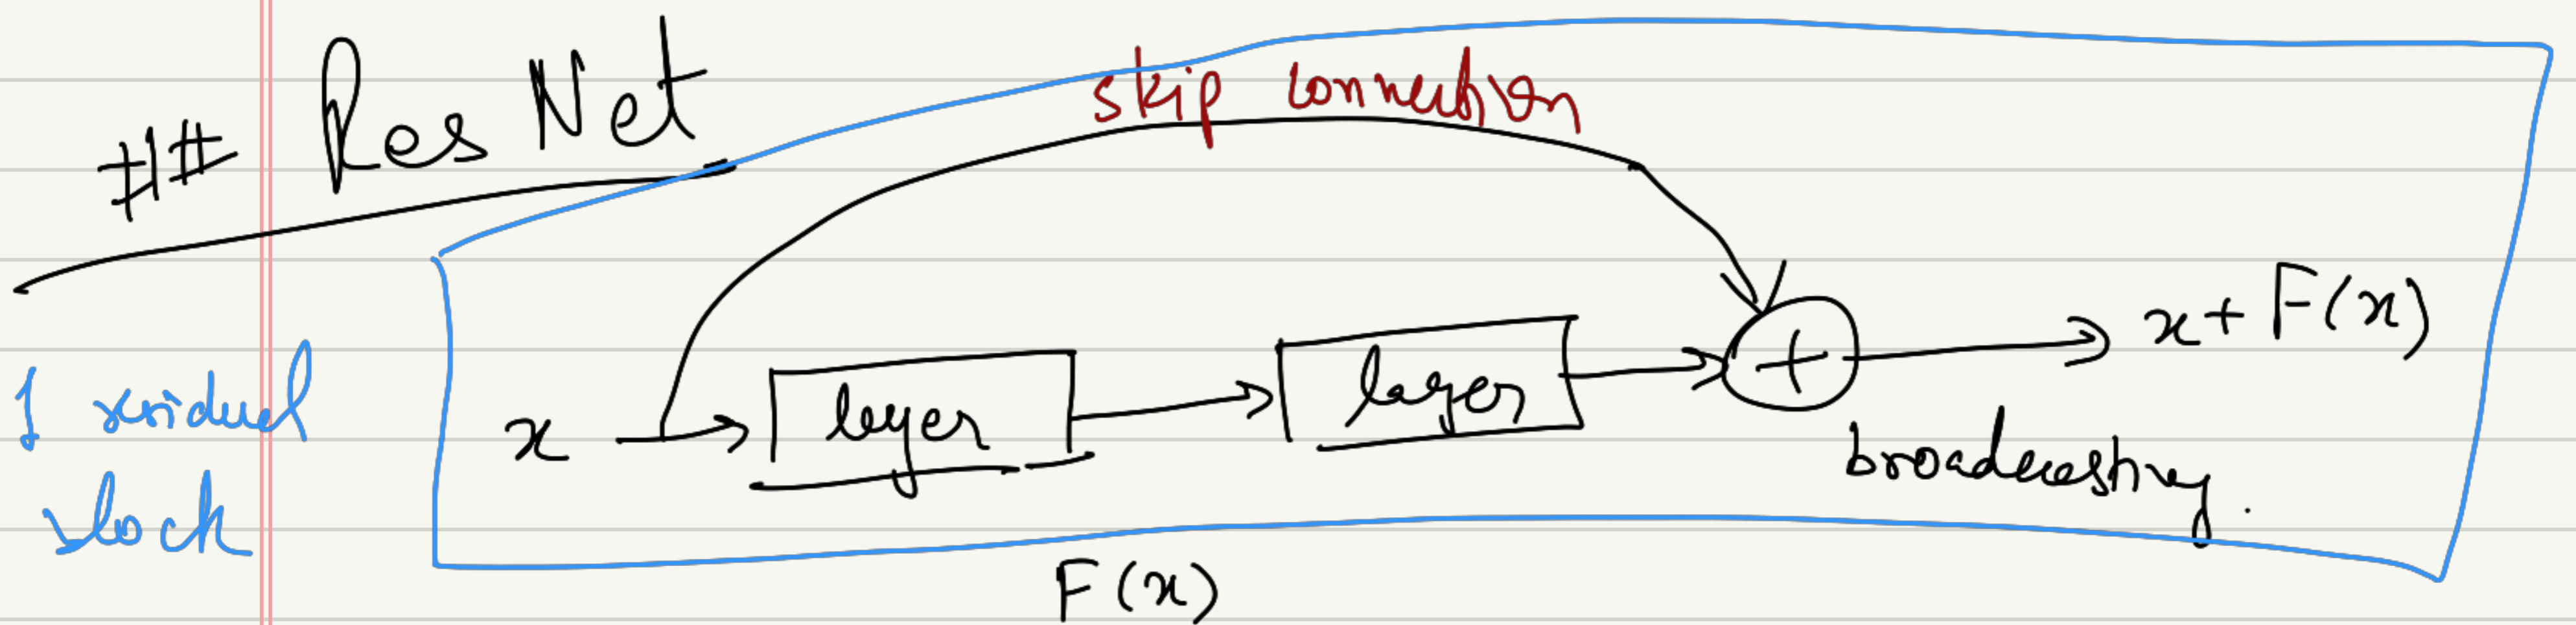
~~backprop~~

$$\frac{\partial L}{\partial W} = \frac{\partial L}{\partial z_3} \cdot \frac{\partial z_3}{\partial z_2} \cdot \frac{\partial z_2}{\partial z_1} \cdot \frac{\partial z_1}{\partial W}$$

back propagation

$$\frac{\partial J}{\partial b} = \left( \downarrow \right) \frac{\partial z_i}{\partial b}$$

• ReLU activation fn  $f(x) = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{if } x \leq 0 \end{cases}$



→ proposed to solve the problem of vanishing / exploding gradients

ResNet is made up of multiple residual blocks.

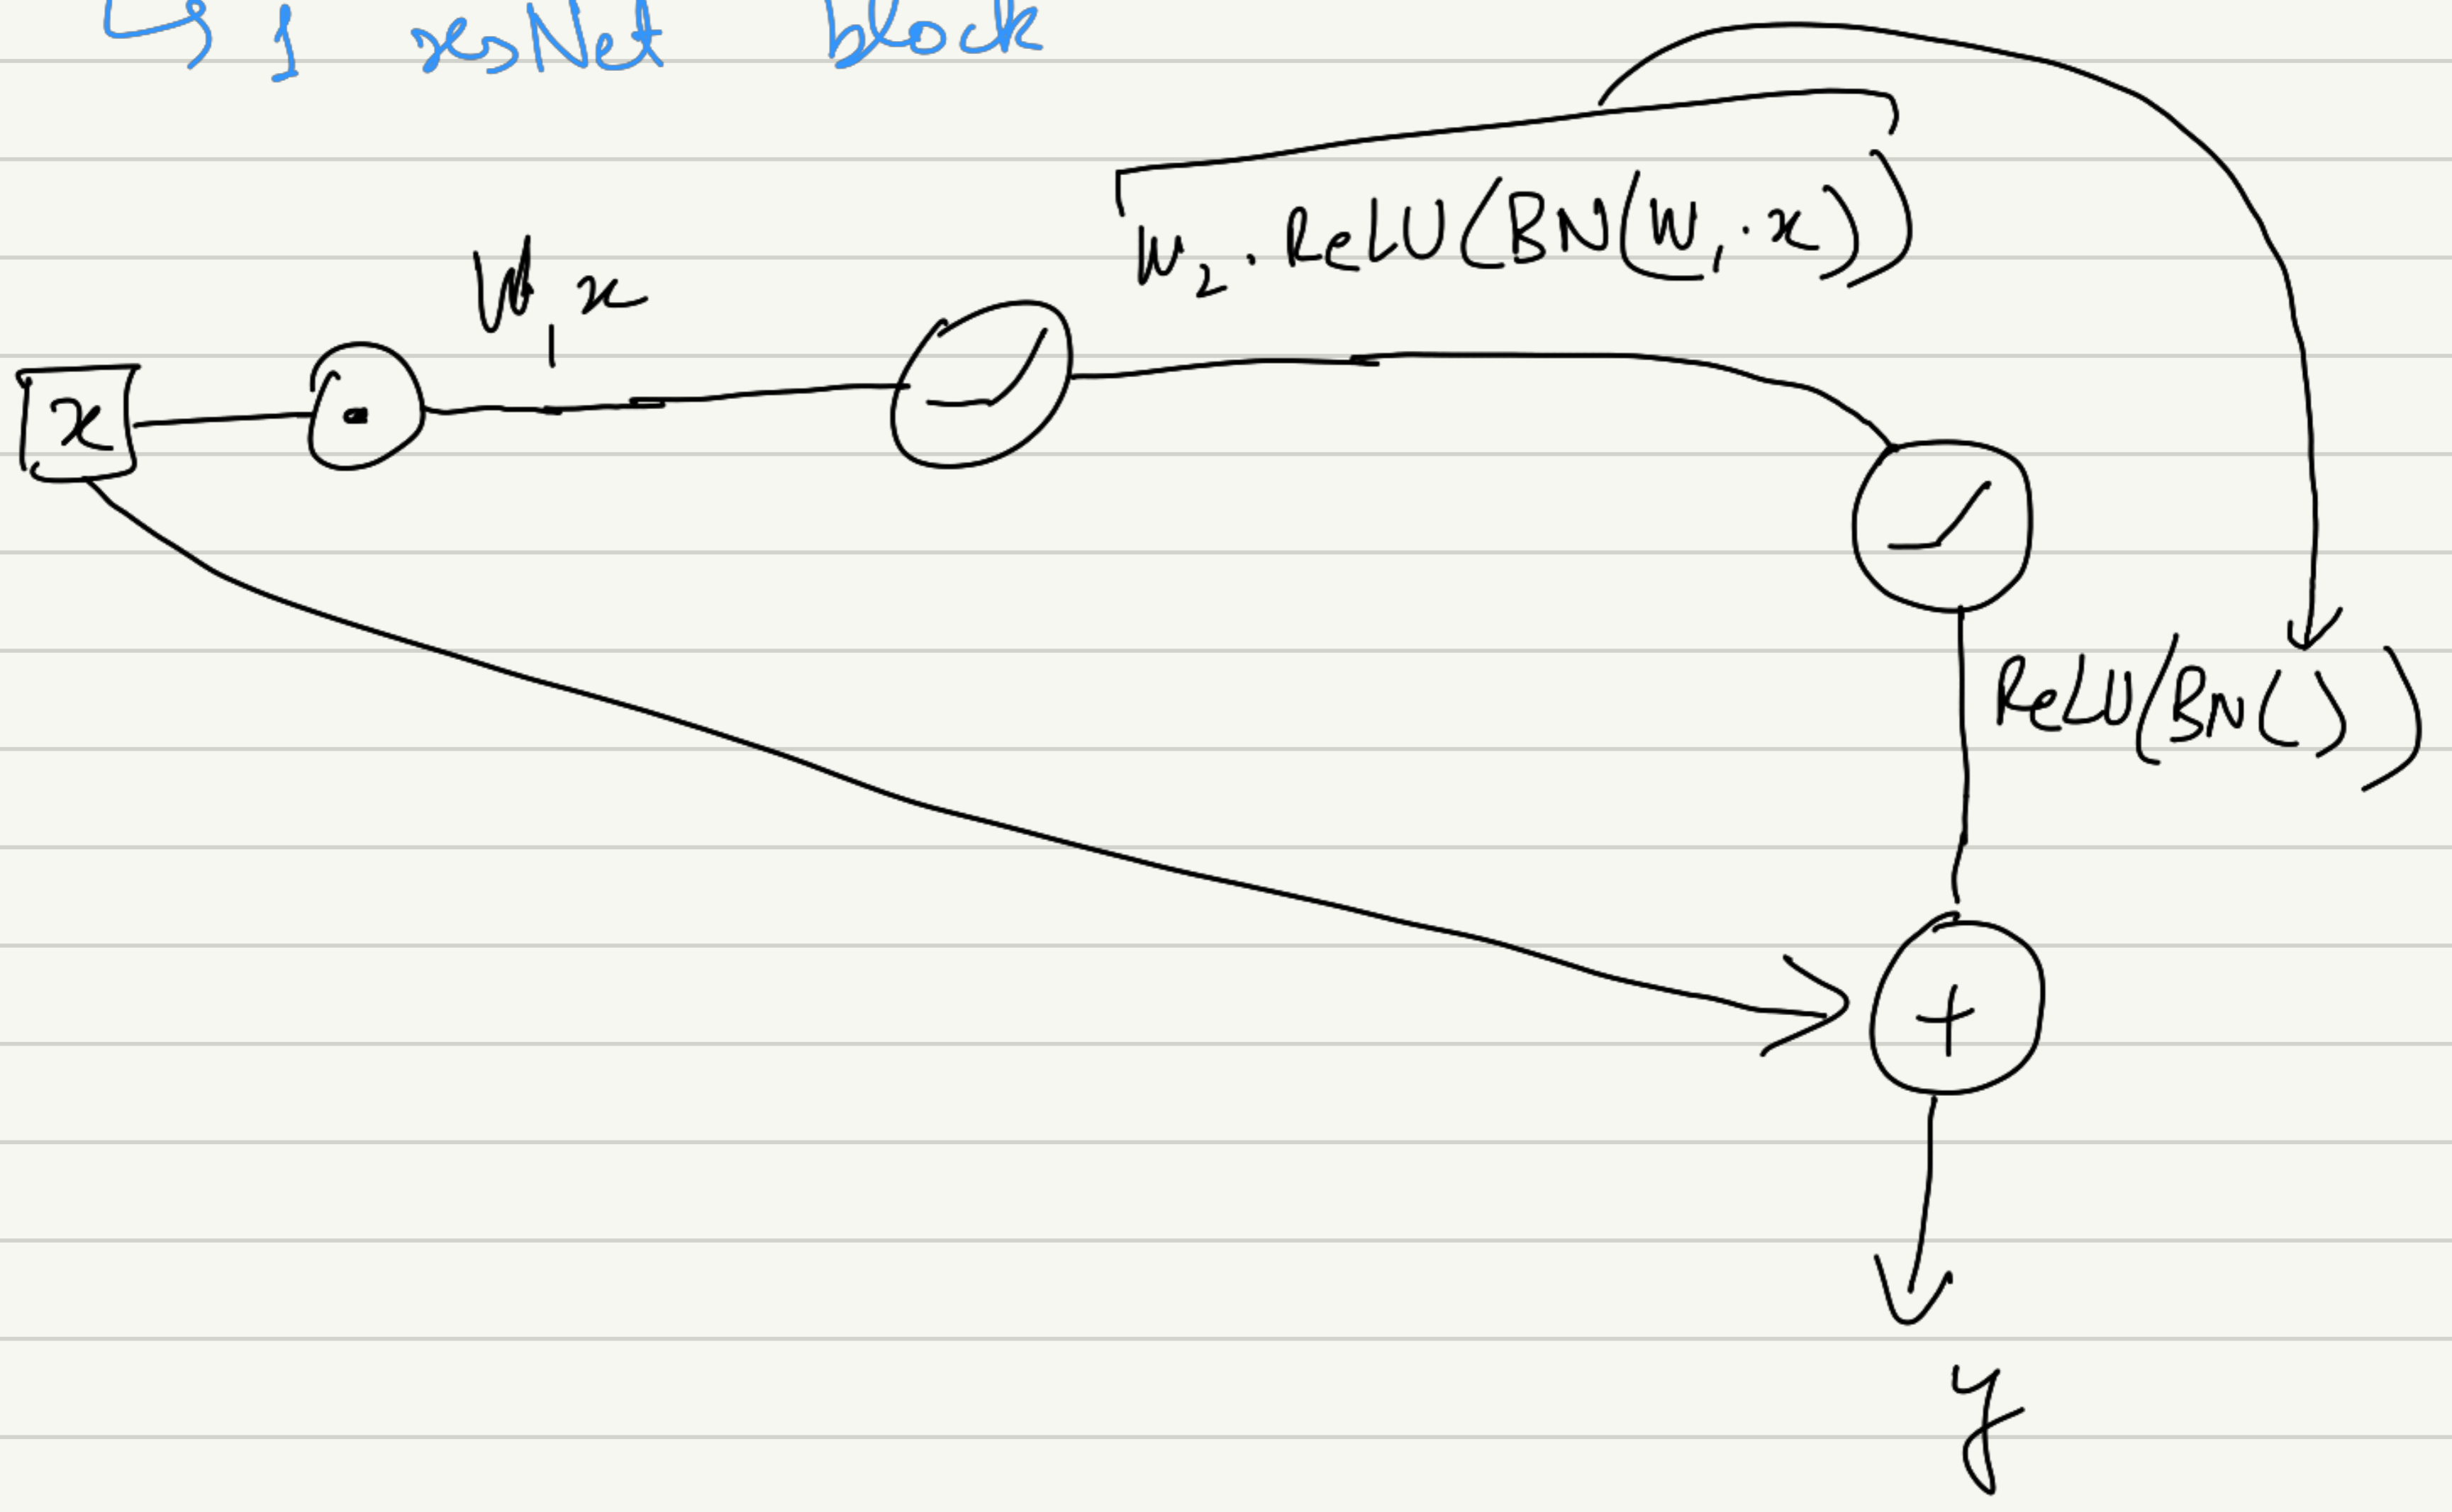
→ regularization??

~~###~~ ~~###~~ ASE - ResNet (Laydi et al. 2025)

Encoding path

$$y = \text{ReLU}(\text{BN}(w_2 \cdot \text{ReLU}(\text{BN}(w_1 \cdot x))) + x)$$

↳ 1 ResNet block



- there are 4 such residual blocks.
- this is standard ResNet, nothing new here

Decoding path

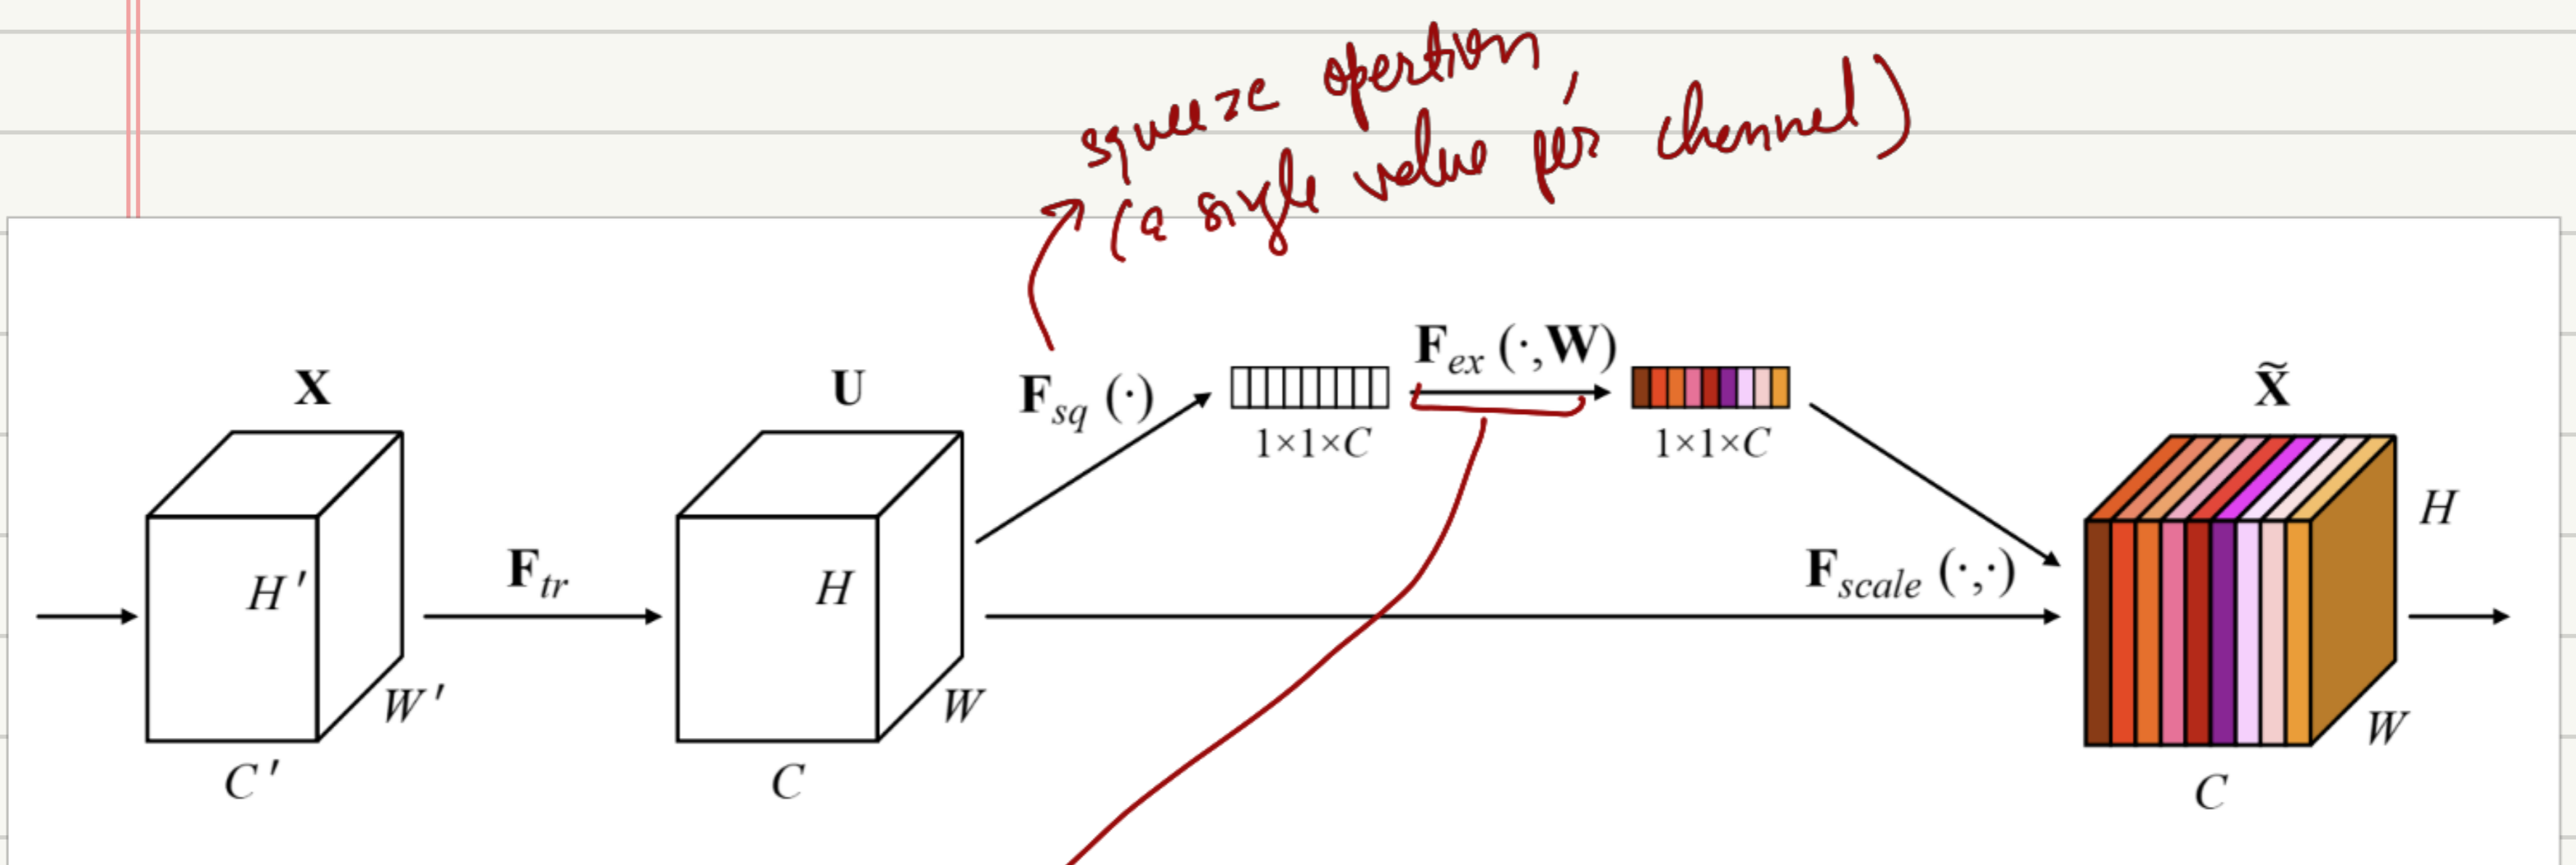
- 4 upsampling blocks.

..... →

# ### Main idea for "Squeeze & excitation"

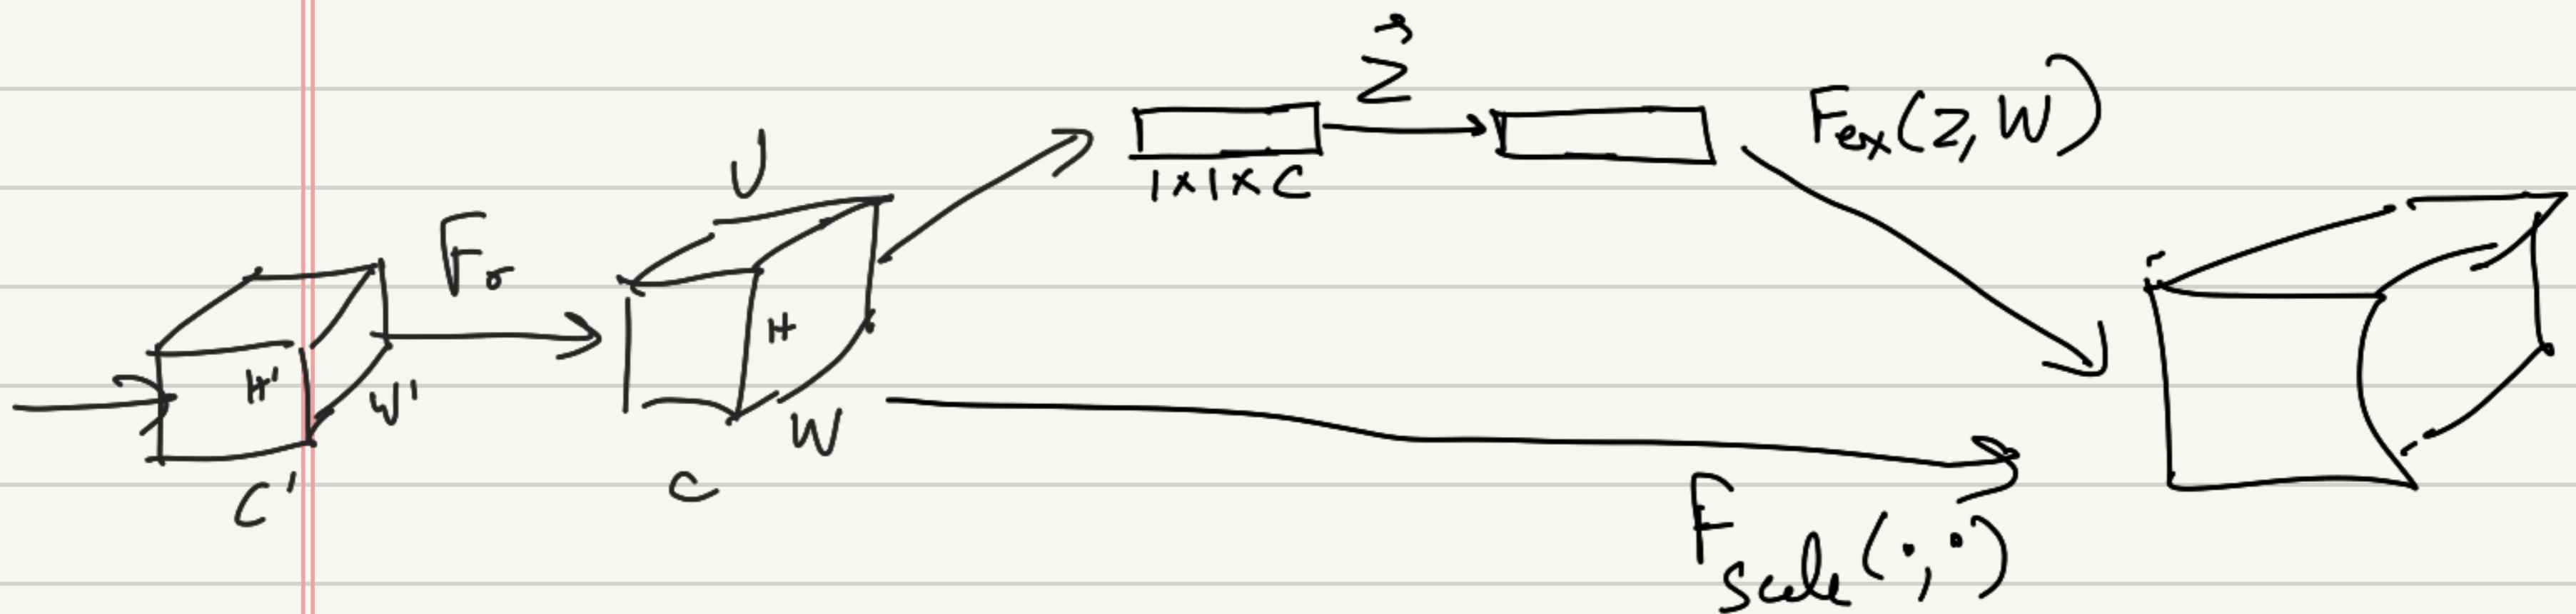
channel wise "attention" mechanism.  
 OR, "attention" is focused on the per channel information.

↳ information in some channels is more important  
 (ex: the 'Green' channel would have more information in an image of the Grass)



excitation operator, to figure out (or assign?) the importance of each channel.

~~(Squeeze & Excitation)~~  
 net  
 (SEnet, Hu et. al. 2019)



## ←...## Decoding path

- noise level:  $N = \sigma(\text{conv2d}(x))$ ,  $N \in \mathbb{R}^{B \times H \times W}$

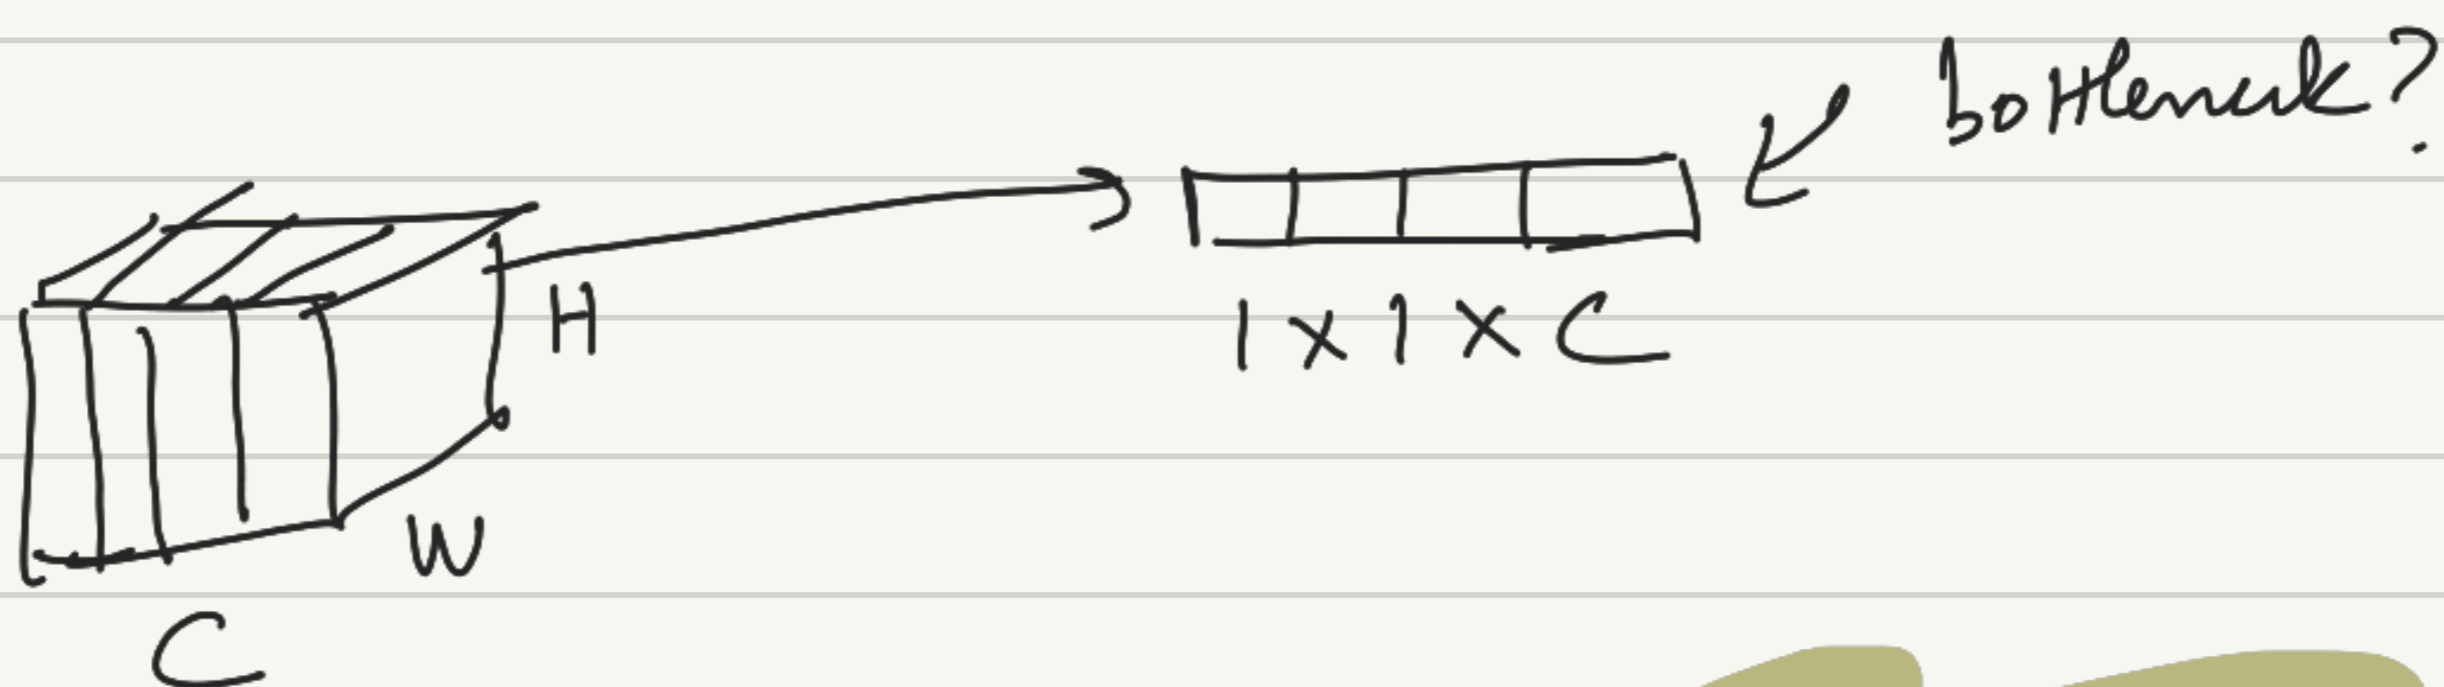
$$\hookrightarrow \text{global average pooling} : \hat{N} = \frac{1}{B \times H \times W} \sum_{b=1}^B \sum_{i=1}^H \sum_{j=1}^W N_{b,i,j}$$

- reductive:  $\max(1, \lceil 16 \cdot (1 + \hat{N}) \rceil)$   
 $\hookrightarrow$  reduction ratio

# Dec 30

Bottleneck & reduction ratio:

In ML, bottleneck is a compression point



Reduction ratio:  $\gamma = \frac{\text{Total channels}}{\text{bottleneck size}}$

eg. ~~for 64 channels in an~~

an image with 64 channels & a reduction ratio 16, would have:

$$\text{Bottleneck size} = \frac{\text{total channels}}{\gamma} = \frac{64}{16} = 4 \text{ neurons}$$

an industry standard to have  $\gamma = 16$  but can be any other value depending on the problem.

\*  
•

regular  
SENet

$$\gamma = 16$$

adaptive :  
SENet

$$\gamma_{\text{adaptive}} = \max(1, 16[1 + \hat{N}])$$

Why does this expression work? What's 'adaptive' about it?

→  $\hat{N}$  (noise lvl) is measured for each channel

↳ low noise ( $\hat{N} \approx 0$ ):  $\gamma \approx 16$  which will lead to standard bottleneck size.

↳ high noise ( $\hat{N} \approx 1$ ):

$$\Rightarrow \gamma = 16(1 + \hat{N}) \approx 16(1 + 1)$$

$$\Rightarrow \gamma \approx 32$$

↳ higher reduction ratio for higher noise level

\*  $\gamma$  is updated on the fly corresponding to the amount of noise per channel.

\* the  $\max(1, \dots)$  is to ensure the reduction ratio is never zero.

... →

### ### Concatenation of feature maps

for two feature maps ~~with~~ with following dimensions:

$$H \times W \times C_1, \quad H \times W \times C_2$$

Concatenation will produce a feature map with dimensions:

$$H \times W \times (C_1 + C_2)$$

for concatenation,

- two out of three dimensions must be the same

- the one that is not the same (usually the channel dimension) is the axis of concatenation.

## loss f<sup>n</sup>s used in leydi et. al. 2025

1.) Multiclass cross entropy loss

from PTM 391 Note 3 supplement:

#### 4. Cross-entropy in information theory/ML

Suppose  $p_i$  = true probability distribution

$q_i$  = predicted probability distribution (e.g. one we derive by a training procedure in ML)

Cross-entropy is:  $H(p, q) = -\sum_i p_i \log_2 q_i$   
(log-loss)

The minimum value of cross-entropy is when  $q_i = p_i$  ("Gibbs' inequality" result)

⇒ minimum value of  $H(p, q)$  is entropy  $H(p) = -\sum_i p_i \log_2 p_i$

The difference between entropy and cross-entropy is the Kullback-Leibler (KL) divergence

$$D_{KL}(p \parallel q) = H(p, q) - H(p)$$

Thus, using cross-entropy as loss function in ML we try to adjust  $q_i$  close to  $p_i$

Example in classification:

|                                |     |     |     |     |           |      |         |
|--------------------------------|-----|-----|-----|-----|-----------|------|---------|
| <b>True distribution:</b>      | 0%  | 0%  | 0%  | 0%  | 100%      | 0%   | 0%      |
|                                | Cat | Dog | Fox | Cow | Red Panda | Bear | Dolphin |
| <b>Predicted distribution:</b> | 2%  | 30% | 45% | 0%  | 25%       | 5%   | 0%      |

← Note entropy of  $p_i$  distribution is zero!  $p_i$  is either zero or one

**Cross-Entropy Loss:**

$$H(p, q) = -\sum_i p_i \log(q_i) \\ = -\log(0.25) = 1.386$$

Classifier



## Dec 31

# Jan 1

# Jan 2

~~Single Gaussian~~

Single Gaussian Covariance  
matrix

$$\begin{array}{c} A \\ u \\ \sigma \end{array} \left[ \begin{array}{ccc} \text{var}(A) & & \\ & \text{var}(u) & \\ & & \text{var}(\sigma) \end{array} \right]$$

Double Gaussian covariance  
matrix

$$\begin{array}{c} A \\ u_1 \\ \sigma_1 \\ B \\ u_2 \\ \sigma_2 \end{array} \left[ \begin{array}{cccccc} \text{var}(A) & & & \text{cov}(A,B) & & \\ & & & & & \\ & & & & & \\ \text{cov}(A,B) & & & \text{var}(B) & & \\ & & & & & \\ & & & & & \end{array} \right] \begin{array}{c} 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array}$$

• by default, scipy uses Levenberg-Marquardt algorithm when

bounds are not specified (source: [https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.curve\\_fit.html#:~:text=curve\\_fit%20is%20for%20local%20optimization,%2Db%20\\*\(%20x\)%20+%20c](https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.curve_fit.html#:~:text=curve_fit%20is%20for%20local%20optimization,%2Db%20*(%20x)%20+%20c))

Gradient descent  
when far from  
optimal solution

Gauss-Newton when  
close to the minimum

— x —

# Jan 3 ### revisiting thermodynamic potentials & phase transitions.

Phase transitions



\* 1<sup>st</sup> order transition  $\Rightarrow$  a discontinuous transition.

Gibbs Enthalpy

$$G = U(S, V, N) - TS + PV$$

$$dG = -SdT + VdP + \mu dN$$

$$\left(\frac{\partial G}{\partial T}\right)_{P, N} = -S + 0 + 0$$

$$V = \left(\frac{\partial G}{\partial P}\right)_T$$

# Jan 4

# Jan 5

# Jan 6

# Jan - 7

1 particle ~~in 1-D~~ in 1-D  $\rightarrow$  state.size() = 3

$$\text{state}[0] = x$$

$$\text{rate}[0] = \dot{x}$$

$$\text{state}[1] = v_x$$

$$\text{rate}[1] = \dot{v}_x$$

$$\text{state}[2] = t$$

$$\text{rate}[2] = \dot{t} = \frac{dt}{dt} = 1$$

this is useful to code system of ODEs to time evolve the system. For example:

~~Velocity verlet~~  
~~also~~  
$$x_{n+1} = x_n + v_n \Delta t + \frac{1}{2} a_n (\Delta t)^2$$

$$v_{n+1} = v_n + \frac{1}{2} (a_{n+1} + a_n) \Delta t$$

to code these, we can write:

$$\text{state}[i] = \text{state}[i] + \text{state}[2 \times i] * dt$$

$$+ \frac{1}{2} \text{rate}[2 \times i] * (dt)^2$$

$$\text{state}[2 \times i] = \text{state}[2 \times i] + \left(\frac{1}{2}\right) ($$

the EOM of a pendulum is given by:

$$\frac{d^2 \theta}{dt^2} = -\frac{g}{L} \sin \theta$$

also,  $v = L \frac{d\theta}{dt}$ ,  $a = \frac{d^2 \theta}{dt^2} = \frac{dv}{dt} = -\frac{g}{L} \sin \theta$

$$\text{state}[0] \rightarrow \theta$$

$$\text{rate}[0] \rightarrow \frac{d\theta}{dt} = v = \text{state}[1]$$

$$\text{state}[1] \rightarrow v$$

$$\text{rate}[1] \rightarrow \frac{dv}{dt} = a = -\frac{g}{L} \sin(\text{state}[0])$$

$$\text{state}[2] \rightarrow t$$

$$\text{rate}[2] \rightarrow \frac{dt}{dt} = 1$$

For a Bouncing Ball

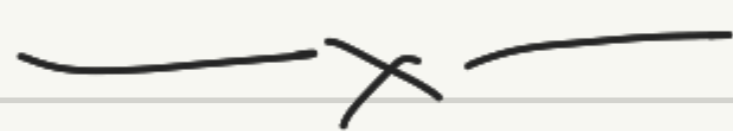
EOM:

$$ma = -mg$$

$$\Rightarrow a = -g$$



~~state [0] → y~~    state [0] → y    rate [0] →  $\frac{dy}{dt} = \text{state [1]}$   
~~state [1] → v~~    state [1] →  $\frac{dy}{dt} \rightarrow v$     rate [1] →  $\frac{dv}{dt} = a = -g$   
state [2] → t    rate [2] → 1



## ## Legendre transforms

•  $f^n$  of a variable  $\xrightarrow{\text{Legendre transform}}$   $f^n$  of conjugate variable

⇒ both  $f^n$ s have same units

e.g. 1    L & H both have units of energy

e.g. 2 :    U, H, F and G (thermodynamic potentials) have units of energy.

(here the conjugate pairs of variables are pressure  $P$ , volume  $V$ , temperature  $T$  & entropy  $S$ )

~~\* thermodynamic variables conjugate~~

• How are conjugate variables defined?

→ from 'Uncertainty principle':  
 $\Delta x \Delta p \geq \frac{h}{2}$

∴  $x$  &  $p$  are conjugate variable pairs

\* Mathematically,  $x$  &  $p$  follow Uncertainty principle b/c they are conjugate variables!!

→ In thermodynamics, the product of conjugate variables have units of energy.

↳  $Work = force \times displacement$

↓  
this holds for any generalized force  
(Pressure & temp. in thermodynamics) &  
generalized displacement (like volume &  
entropy in thermodynamics)

\* So as a rule of thumb, a pair of generalized force & generalized velocity makes a conjugate variable pair.

D loop of actin monomer

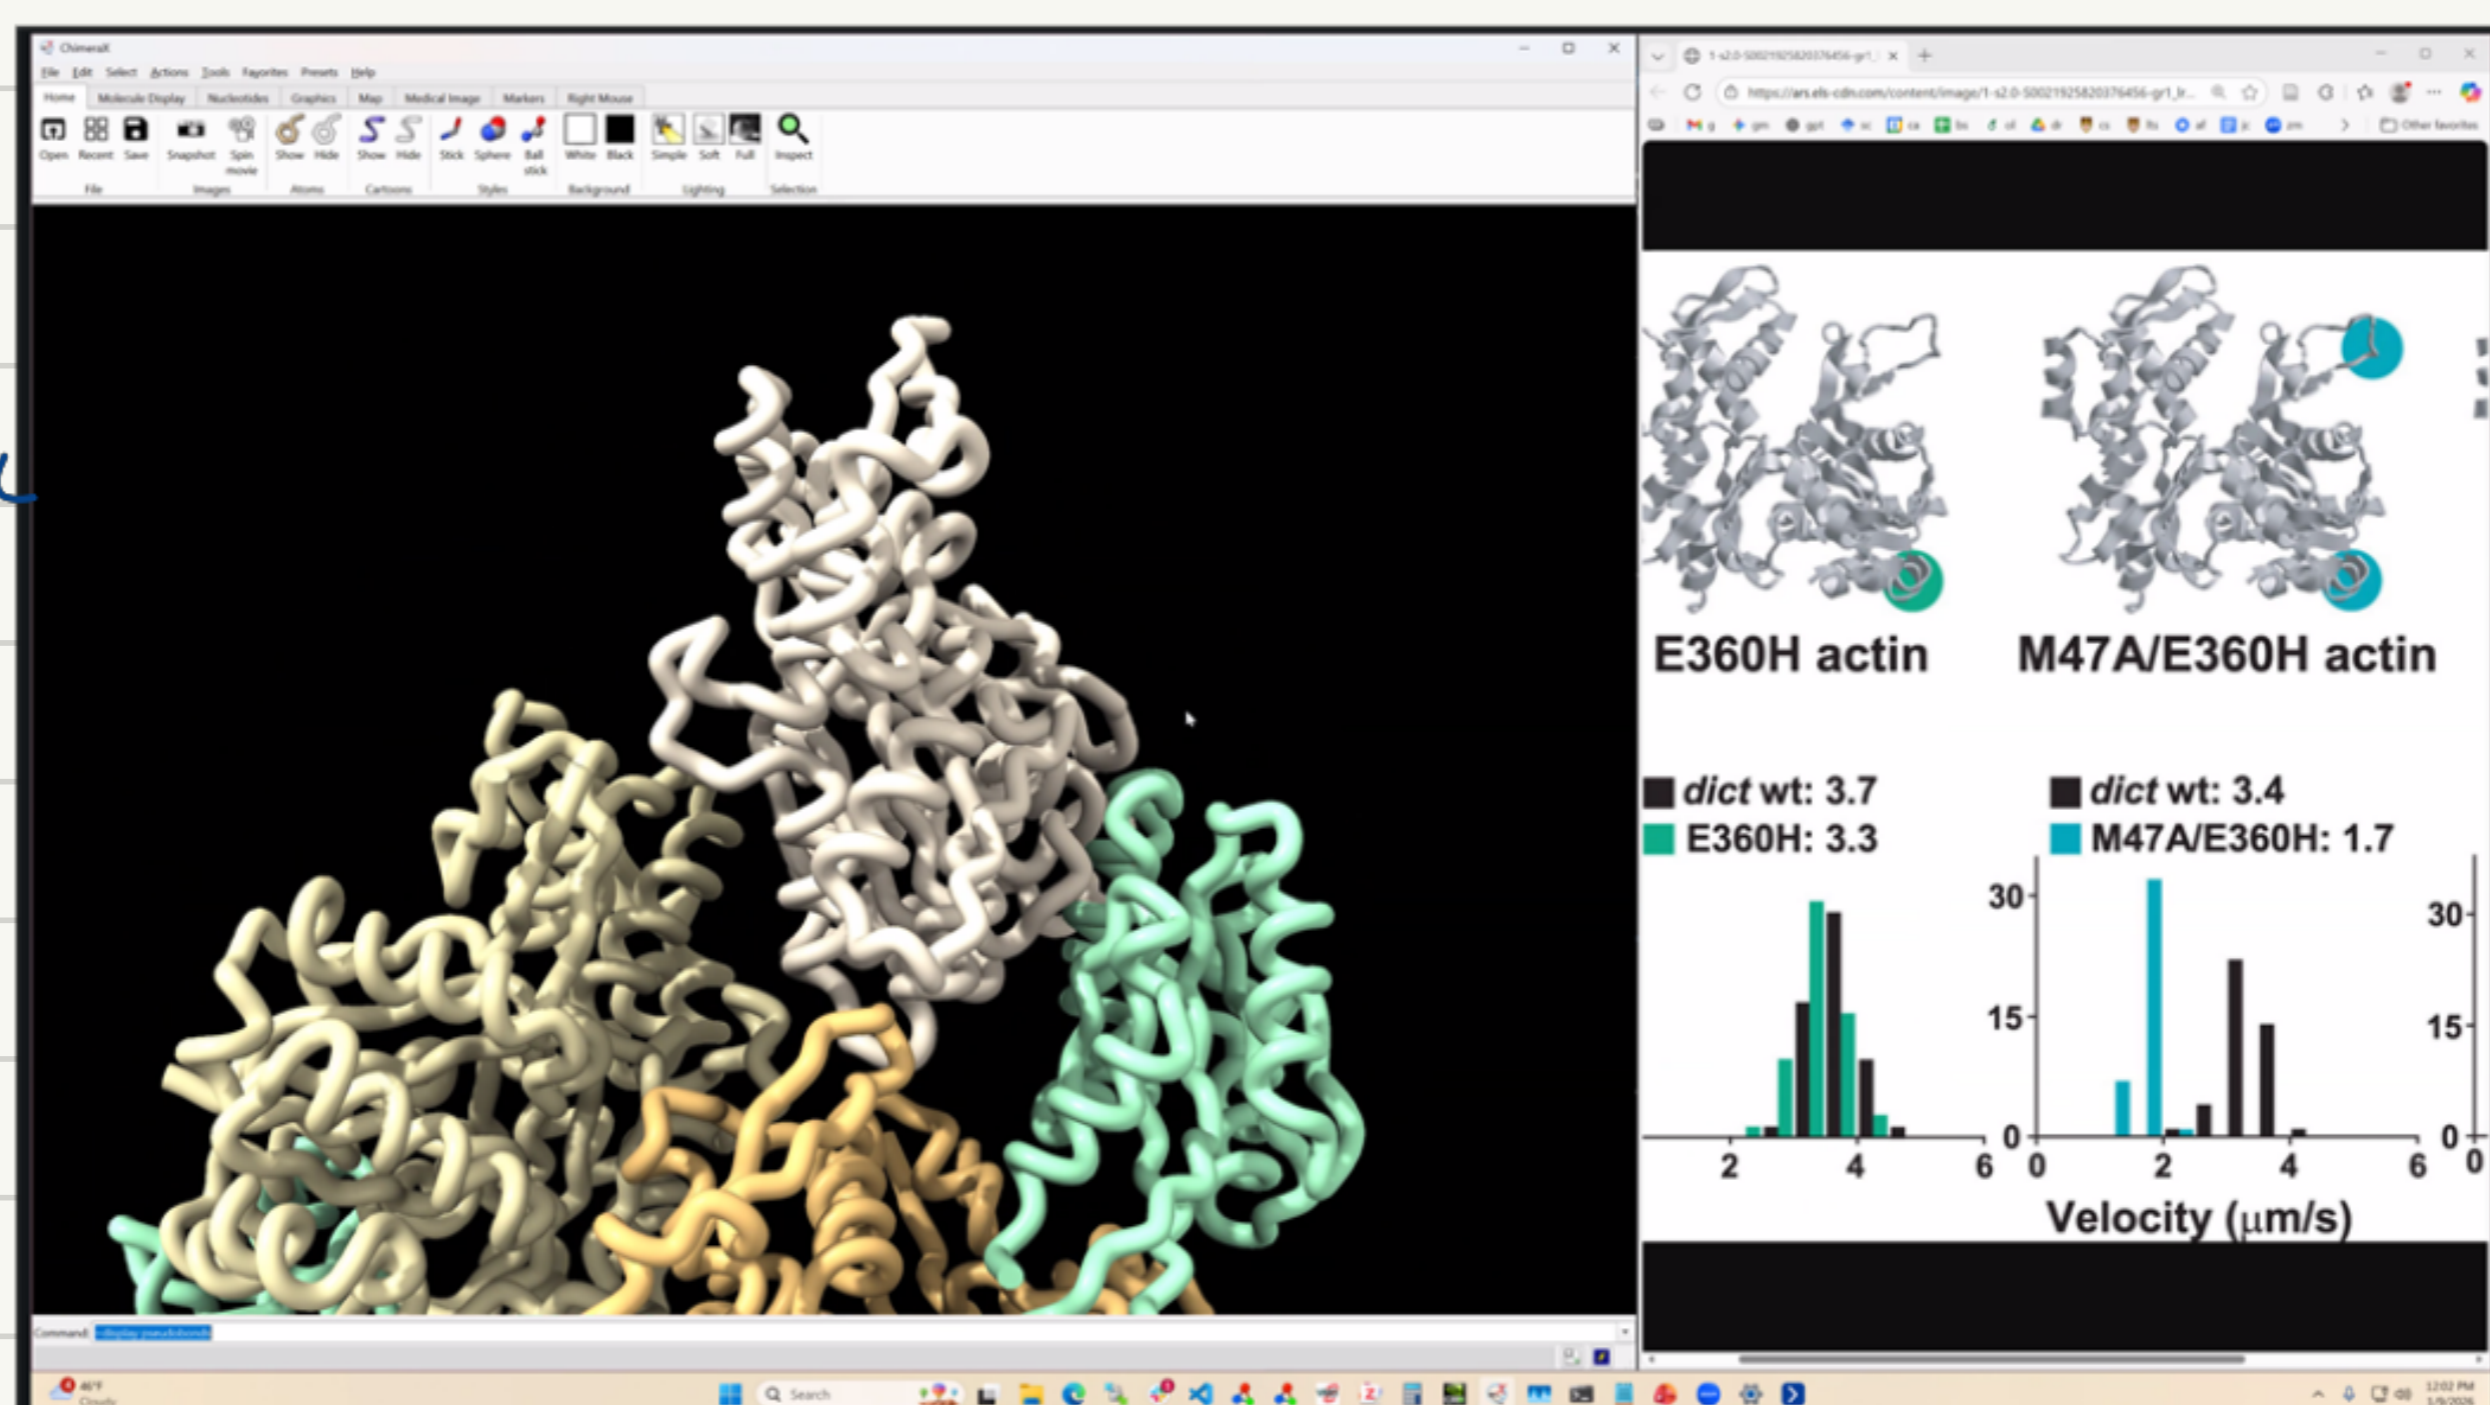
↳ cofilin breaks the D-loop connection. on ~~two~~

↳ on two actins.

\* site of hydrolysis on actin.

↳ monomeric actin structure.

\* groove portion  
of the hydrophobic  
plegue



→ thronin

↳ another severing protein like cofilin

→ ~~pointed~~ pointed end polymerization.

↳ hydrolysis

↳ detailed balance

↳ prof. Vanyukov thought about  
this for a year.

= plot counts histogram. } over time  
↳ most legitime.

→ log-log plot.  
↳ to include the outliers

→ ~~test~~ remove spurious and analyze  
↳ don't have to keep verifying  
↳ another graphs.

↳ bias vs. proportion etc.

→ ~~Summary~~ for Dipt

↳ get back to Bond