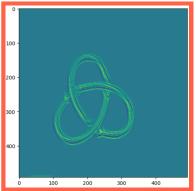
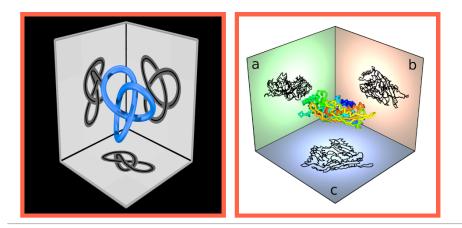
Picture recognition, knots and quantum algebra

Or: Knot detector



 $\ensuremath{\mathsf{Accept}}\ensuremath{\mathsf{Change}}$ what you cannot $\ensuremath{\mathsf{changeaccept}}$

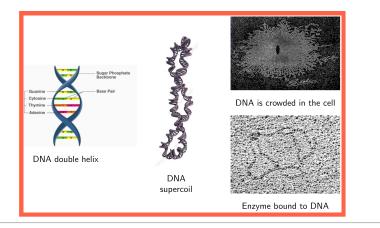
I report on work of Dranowski-Kabkov-Melamud (via Yulia's dream)



• Knot = closed string (a circle S^1) in three spaces; link = multiple components

Knots are studied by projections to the plane Shadows

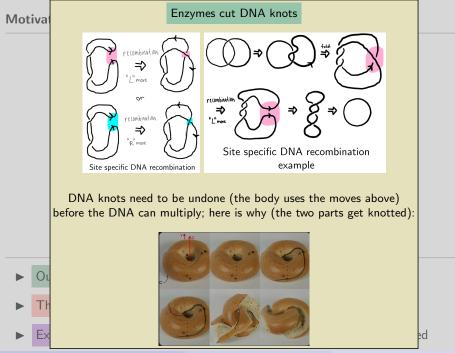
In math knot theory started in the early 20th century

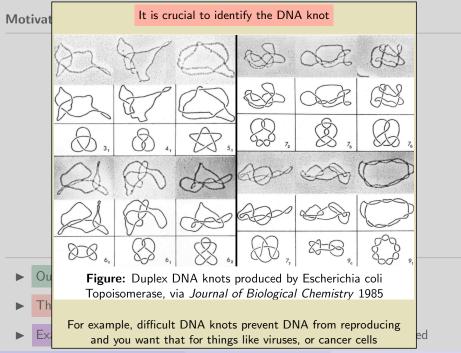


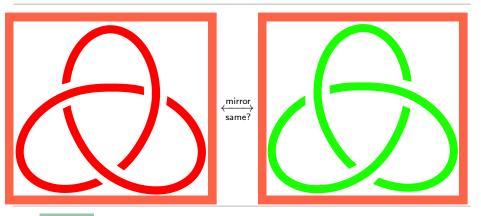
► Outside math knot theory started ~1980

Theorem Long strings in small boxes get knotted

• Example DNA is a long string in a tiny box (cell) \Rightarrow its often knotted

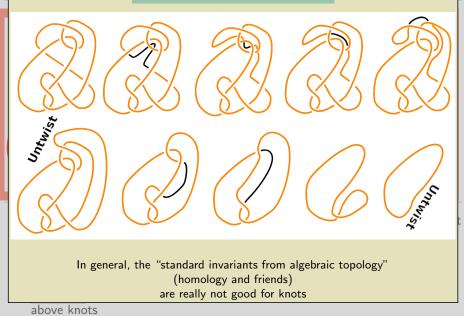


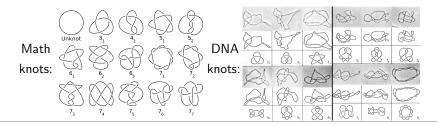




- Problem Deciding whether two knot projections are the same knot is difficult
- Task Find an invariant. Sounds easy? Well, most knot invariants are pretty bad...so: find a 'good' knot invariant
- Example There was no knot invariant before 1980 that can distinguish the above knots

Even the unknotting problem is tricky





- ▶ Knots are hard for humans/machines to recognize visually
- Developing an automated method to extract data from knot images has impact beyond mathematics (DNA, protein folding, polymer science, ...)
- ► Goal Develop a CNN model for knot recognition
- ► Aims, explained together with why this is feasible as we go:
 - Identify the number of crossings. (Done!)
 - Produce planar diagram (PD) presentations for computation of invariants, e.g., Jones polynomial. (On it!)

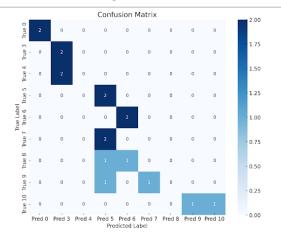
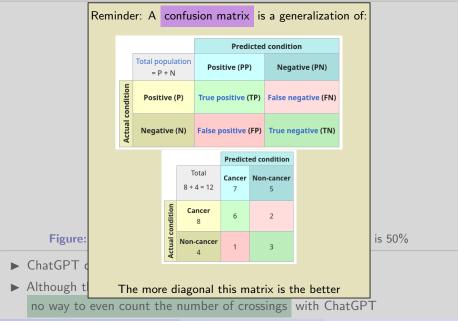


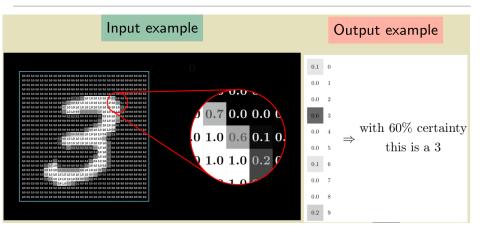
Figure: Confusion matrix of o4-mini results. Average accuracy is 50%

- ► ChatGPT did poorly
- ► Although the test set was small, there was

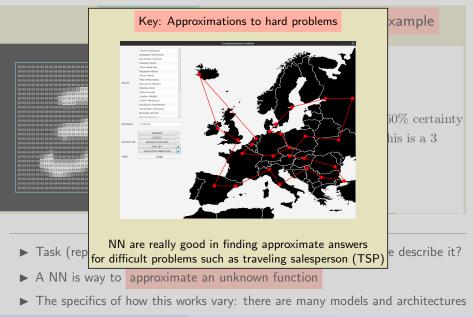
no way to even count the number of crossings with ChatGPT

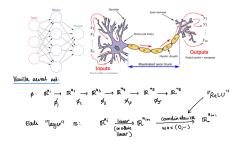


- ► Task: Identify handwritten digits
- ▶ We can see this as a function in the following way:
 - (a) Convert the pictures into grayscale values, e.g. 28×28 grid of numbers
 - (b) Flatten this into a vector, e.g. $28 \times 28 \mapsto$ a vector with $28^2 = 784$ entries
 - (c) The output is a vector with 10 entries
 - (d) We thus have a function $\mathbb{R}^{784} \to \mathbb{R}^{10}$



- ▶ Task (rephrased): We have a function $\mathbb{R}^{784} \to \mathbb{R}^{10}$. How can we describe it?
- ► A NN is way to approximate an unknown function
- ▶ The specifics of how this works vary: there are many models and architectures





A typical model architecture consists of a composition of functions, say $\phi_1, \phi_2, ..., \phi_N$. This composition is applied to the input data, and the result is processed to optimize the model parameters called weights

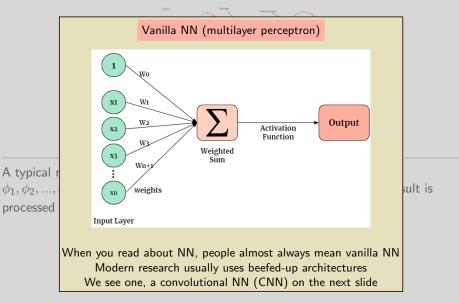
$$\phi_{\mathsf{N}} \circ \phi_{\mathsf{N}-1} \circ ... \circ \phi_1(\mathsf{input}) = \mathsf{output}$$

$$\phi(\mathbf{x};\boldsymbol{\theta}) = \phi_{N} \circ \phi_{N-1} \circ \dots \circ \phi_{1}(\mathbf{x}),$$

where **x** is the input and θ is the vector of all weights

Picture recognition, knots and quantum algebra

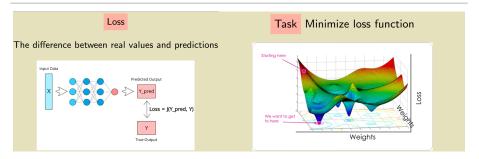
Or: Knot detector



where **x** is the input and θ is the vector of all weights

Picture recognition, knots and quantum algebra

Or: Knot detector

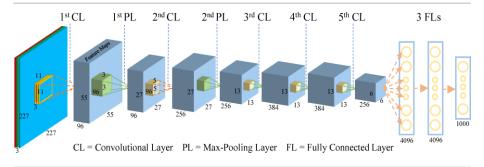


Loss function *L* measures the discrepancy between predictions and ground truth:

$$L = L(predictions, truths) = L(\phi(input), truths) = discrepancy$$

$$\min_{\boldsymbol{\theta} \in \text{dom } L} L \quad \Rightarrow \quad \boldsymbol{\theta}^{(l+1)} = \boldsymbol{\theta}^{(l)} - \eta \nabla L(\boldsymbol{\theta}^{(l)}), \tag{1}$$

where the superscript (1) denotes the *I*-th iteration, η is the learning rate, and ∇L is the gradient of the loss function with respect to the weights



- Task Classify images by detecting local patterns
- Key idea Convolutions detect features (edges, textures)
- CNN structure:
 - (a) Convolutional layers: Extract features
 - (b) Pooling layers: Downsample the image Ignore for today
 - (c) Fully connected layers: Classify the image Ignore for today

Let us open https://poloclub.github.io/cnn-explainer/

Pre

Idea (kernel, filter, etc.)

Operation	Kernel ω	Image result g(x,y)
Identity	$\left[\begin{array}{rrrr} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{array} \right]$	
Ridge or edge detection	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$	
Sharpen	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
Box blur (normalized)	$\frac{1}{9} \left[\begin{array}{rrrr} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{array} \right]$	C.

In image processing, a kernel, convolution matrix, or mask is a small matrix used for blurring, sharpening, embossing, edge detection, and more This is accomplished by doing a convolution between the kernel and an image

Picture recognition, knots and quantum algebra

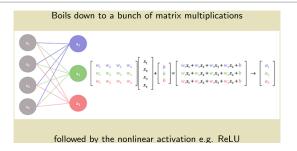
Or: Knot detector



Figure: Steps of image preprocessing

▶ Our data source were many knot images as in (a)

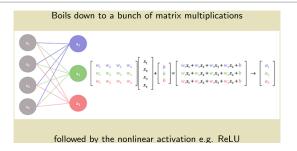
▶ We ran a skeletonization on it



The easiest architecture is linear layer + nonlinear activation function e.g. ELU.

 $FC(\mathbf{x}) = ELU(\mathbf{W}\mathbf{x} + \mathbf{b})$

$$Vanilla = FC^k, \ k \in \mathbb{N}$$

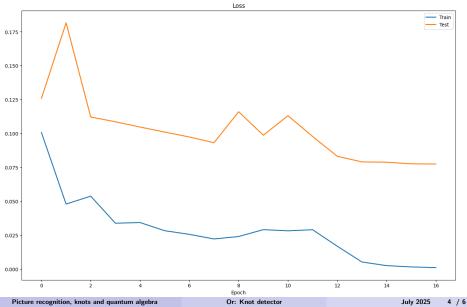


The easiest architecture is linear layer + nonlinear activation function e.g. ELU.

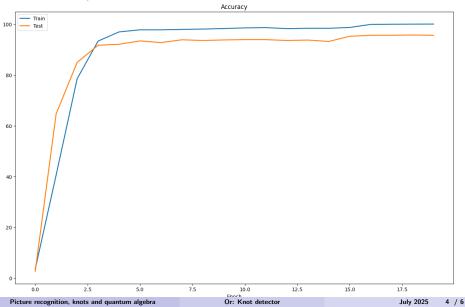
 $FC(\mathbf{x}) = ELU(\mathbf{W}\mathbf{x} + \mathbf{b})$

$$Vanilla = FC^k, \ k \in \mathbb{N}$$

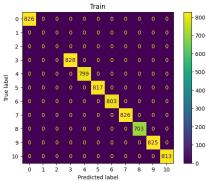


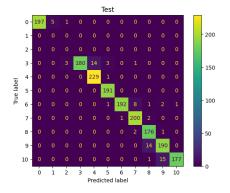


Vanilla, accuracy



Vanilla, confusion matrices





Vanilla, weights

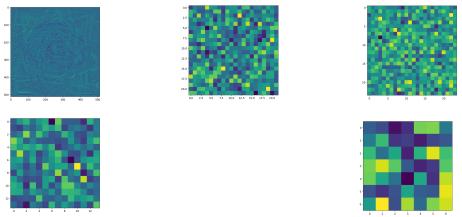
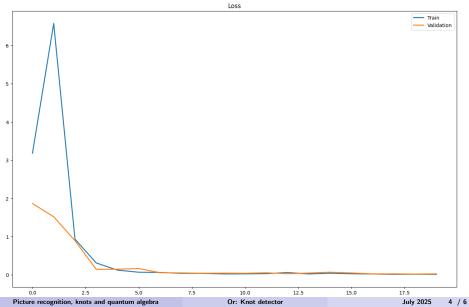


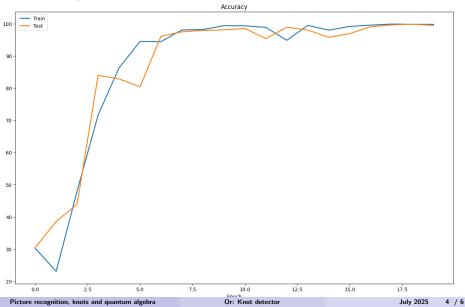
Figure: Weights of different Vanilla's layers - pretty random!

Problem: This is random, doesn't seem generalizable

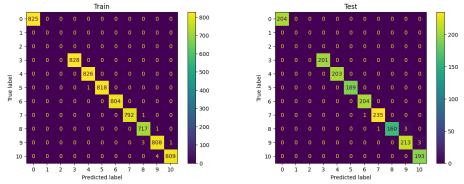
CNN, loss



CNN, accuracy



CNN, confusion matrices



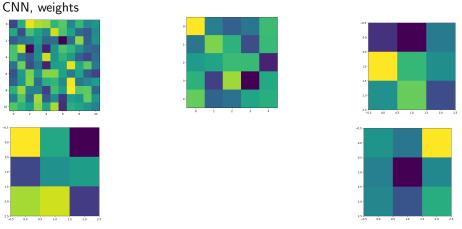
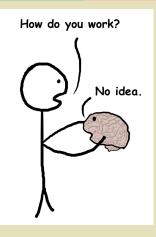


Figure: Weights of different CNN's layers

Turns out this is not random: we feed in a knot next!

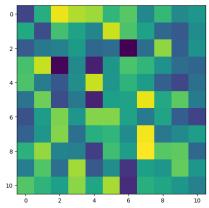
Is a NN a black box? Well...yes and no!

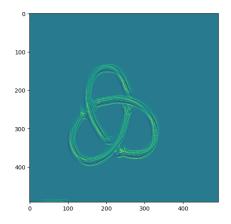


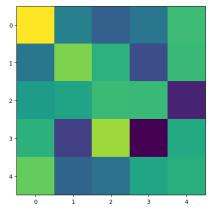
The vanilla NN learns the pixel distribution – and that is why it looks random! Most vanilla NN do something similar – and that is why people call them black boxes For more sophisticated architectures it is often possible to understand what is going on! We see that next for the CNN – it learns what crossings are!

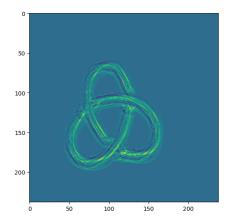
Picture recognition, knots and quantum algebra

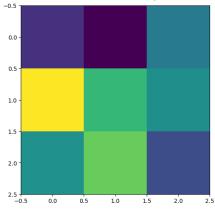
Or: Knot detector

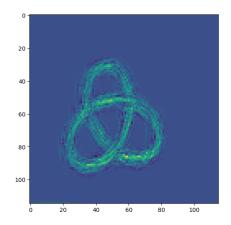


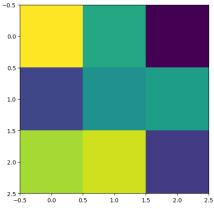


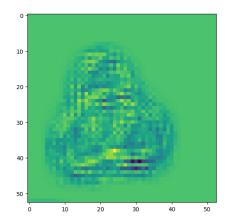


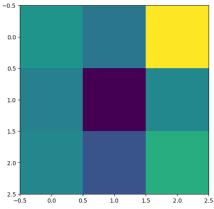


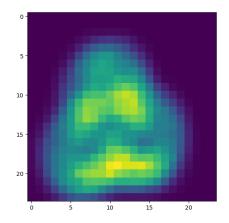




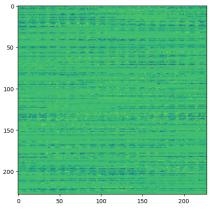


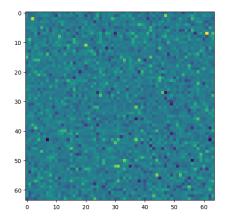




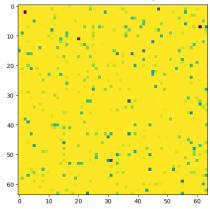


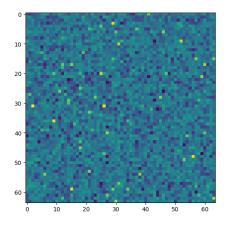
What does the CNN actually do?



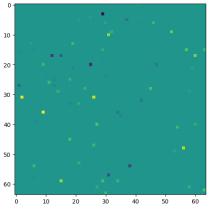


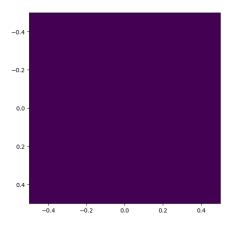
What does the CNN actually do?





What does the CNN actually do?





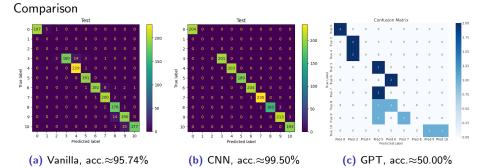
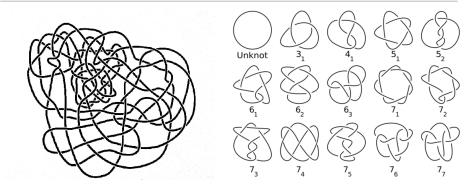


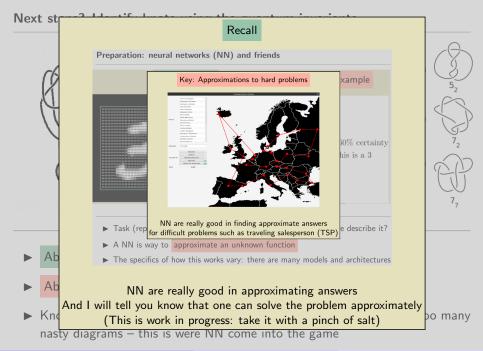
Figure: Confusion matrices on test dataset

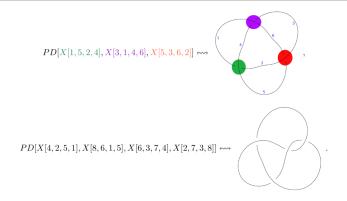


Above, left A nasty picture of the unknot

Above, right The first few knots (easy diagrams)

► Knot detection is not (expected to be) an easy problem: there are too many nasty diagrams – this is were NN come into the game



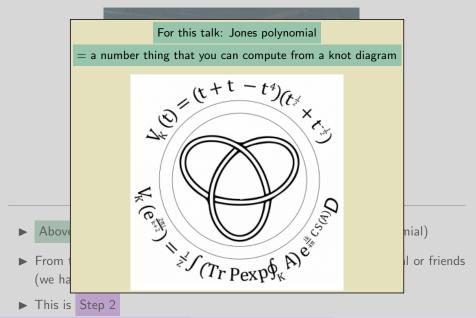


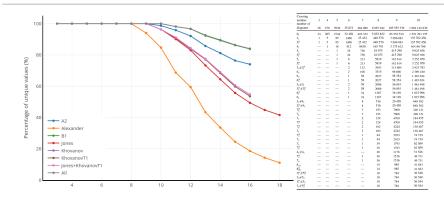
• Above A PD presentation (an algebraic way to encode knot diagrams)

- ▶ Using the same architecture as before, we can read this in from an image
- ► This is Step 1



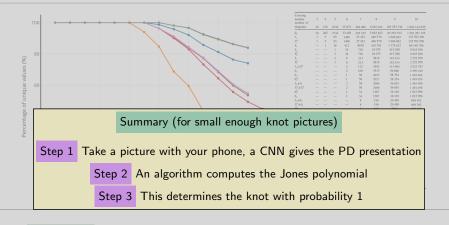
- Above Jones receives the fields medal (also for the Jones polynomial)
- ► From the PD presentation one can compute the Jones polynomial or friends (we have a program that can do this fast for 100 crossing knots)
- ► This is Step 2





"Problem" The Jones polynomial is not a very good invariant

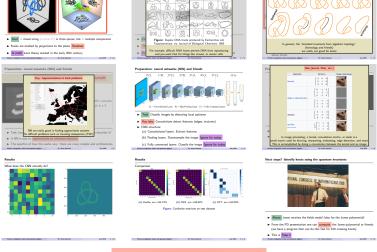
- Solution This doesn't matter as the distribution of knot diagrams is biased towards "easy knots" – and the polynomial is good on these!
- ► This is Step 3



"Problem" The Jones polynomial is not a very good invariant

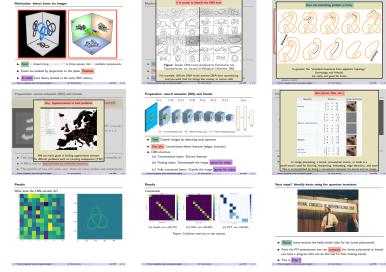
Solution This doesn't matter as the distribution of knot diagrams is biased towards "easy knots" – and the polynomial is good on these!

► This is Step 3



There is still much to do...

Motivation: detect knots via images



Thanks for your attention!