Group Equivariant Non-Expansive Operators:

A blueprint for a different kind of Neural Networks

Hello!

I am Giovanni Bocchi

2^{nd} year PhD student working on applications of Group Equivariant Non-Expansive Operators to Artificial Intelligence in the direction of Explainability an Trustworthiness.

1

Framework

Statistical Learning

The Framework of Statistical Learning

Consider a **data domain $\boldsymbol{\mathcal{X}}$**, a **labels set $\boldsymbol{\mathcal{Y}}$** and a joint unknown **probability distribution $\boldsymbol{\mathcal{D}}$** on $\mathcal{X} \times \mathcal{Y}$.

A function $h: \mathcal{X} \to \mathcal{Y}$ is called **predictor** (**regressor** or **classifier** depending on the problem).

Undesired predictions are penalized through a **loss function**

$$\ell: \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$$

Example: $\ell(y, \widehat{y}) = (y - \widehat{y})^2$

The Framework of Statistical Learning

The objective is to find a predictor that minimizes the **Statistical Risk** among a certain class of predictors:

$$\ell_{\mathcal{D}}(h) = \mathbb{E}_{(X,Y)\sim\mathcal{D}}[\ell(Y, h(X))]$$

The usual strategy is to observe an i.i.d. **sample** (usually called **training set**) $S = (X_i, Y_i)_{i=1}^{n}$ and to try to minimize the **Empirical Risk** instead

$$\hat{\ell}(h) = \frac{1}{n}\sum_{i=1}^{n} \ell(Y_i, h_S(X_i))$$

2

Neural Networks

A popular class of predictors

What are Neural Networks?

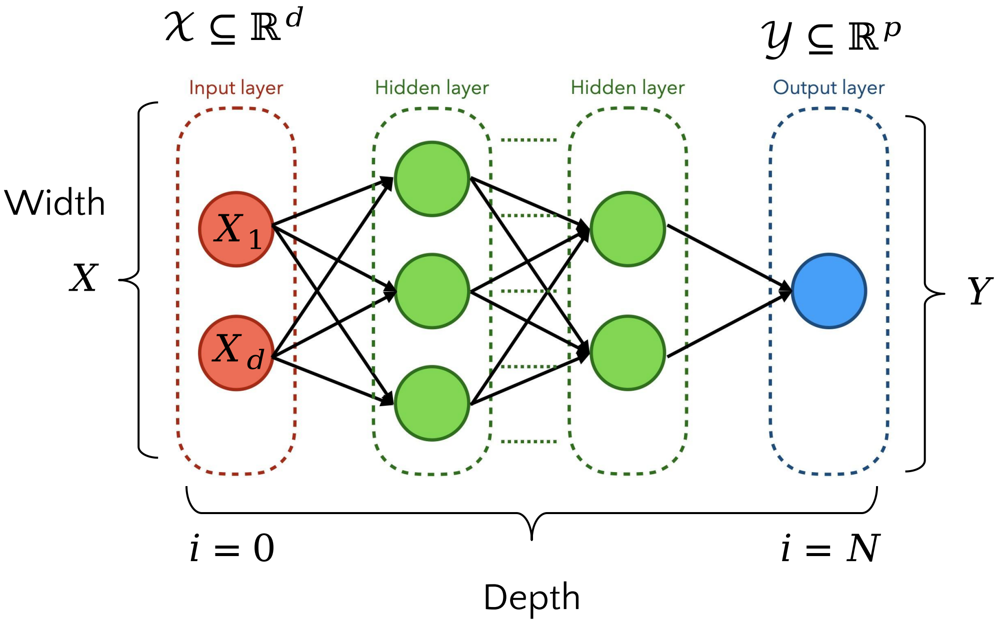

Feedforward NN:

$$H_0 = X$$
$$H_{k+1} = \sigma_k(W_k H_k + b_k)$$
$$H_N = \widehat{Y}$$

$$NN = \phi_N \circ \cdots \circ \phi_0$$
$$\phi_0 = Id$$
$$\phi_k = \sigma_k \circ (W_k \cdot \; + b_k)$$

A *classical* result

Theorem (Cybenko 1989)

Let $C([0,1]^n, \mathbb{R})$ denote the set of all continous function from the unit hypercube to $\mathbb{R}$, let σ be any sigmoidal activation function then the finite sum of the form

$$f(x) = \sum_{i=1}^{N} \alpha_i \sigma(w_i^t x + b_i)$$

is dense in $C([0,1]^n, \mathbb{R})$.

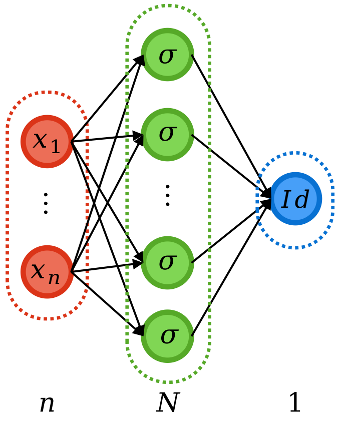

A *recent* result

Theorem (Lu et al. 2017)

For any Lebesgue integrable function $f : \mathbb{R}^n \to \mathbb{R}$ and any $\varepsilon > 0$ there exists a fully-connected ReLU network A with width $d_m \leq n + 4$, such that the function F_A represented by this network satisfies:

$$\|f - F_A\|_{L^1} < \varepsilon$$

Function Approximation

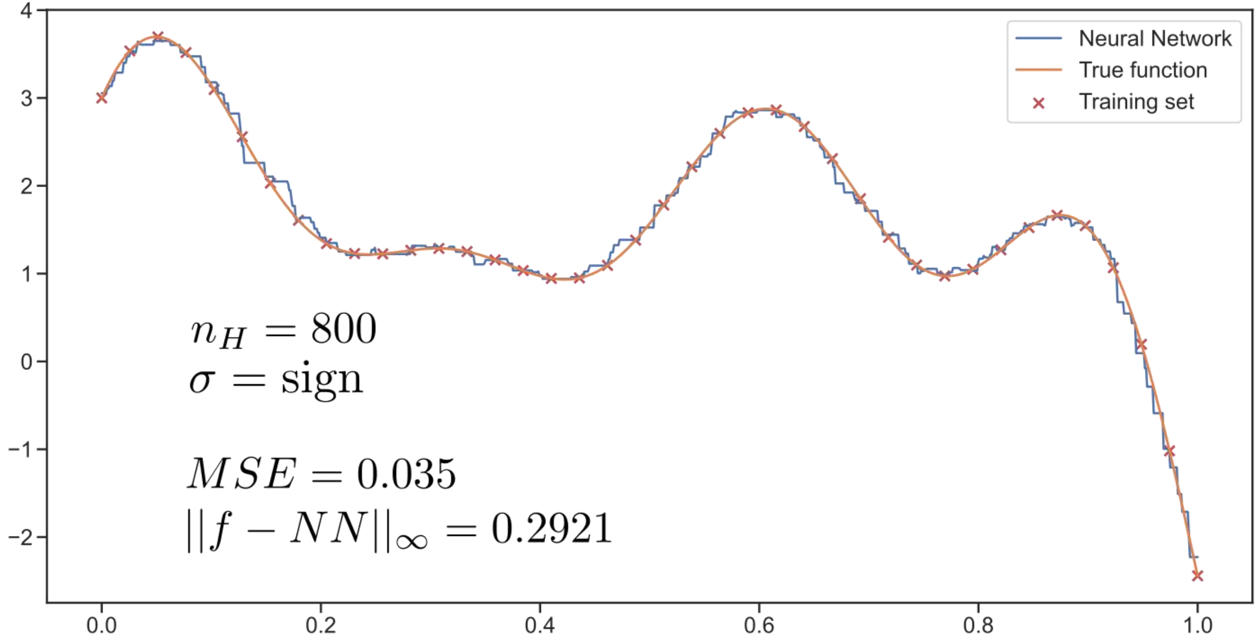

The Zoo of NN

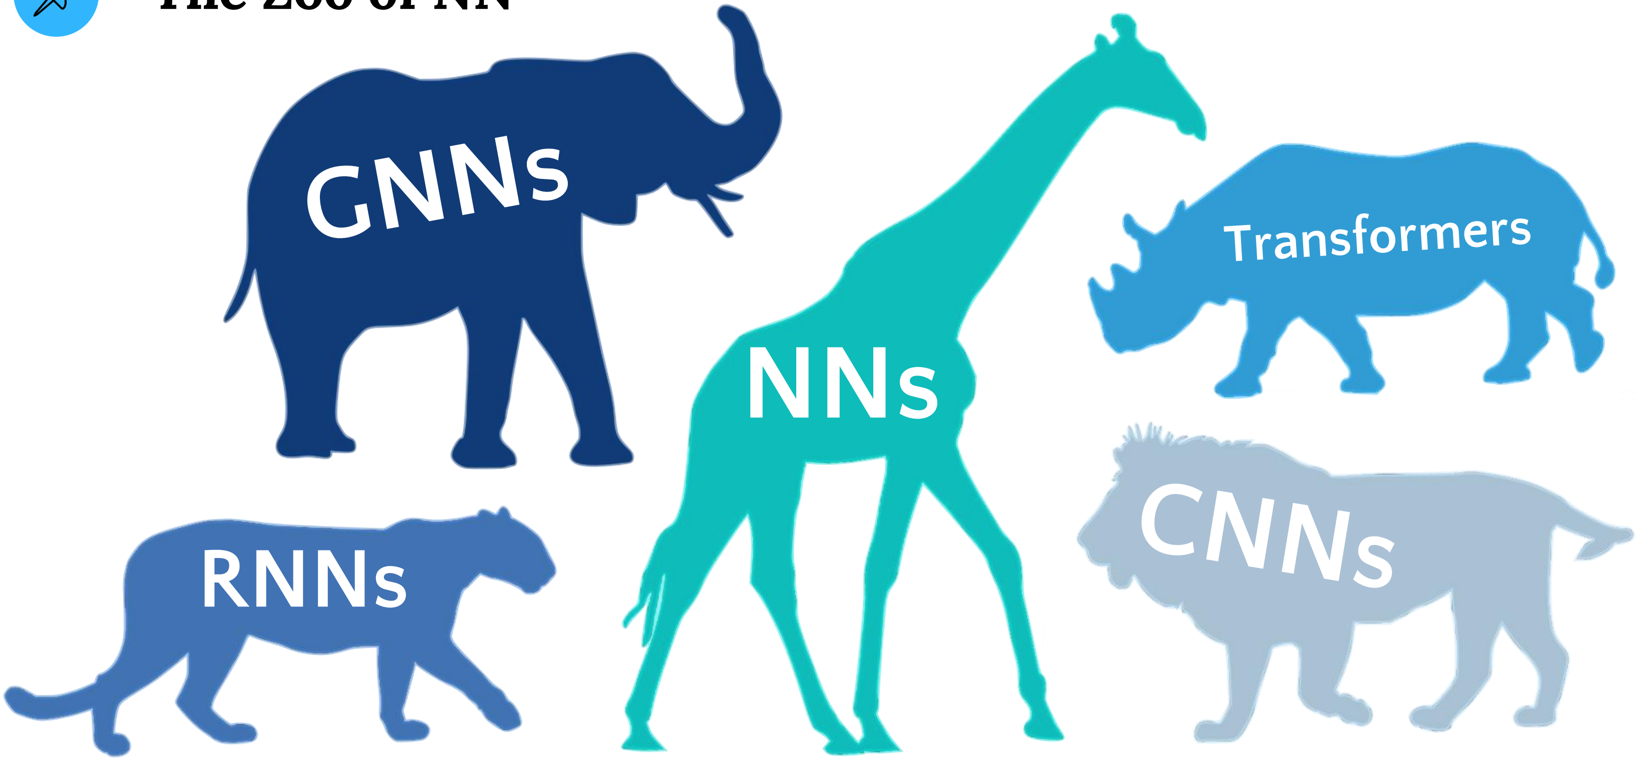

An animal from the Zoo: CNNs

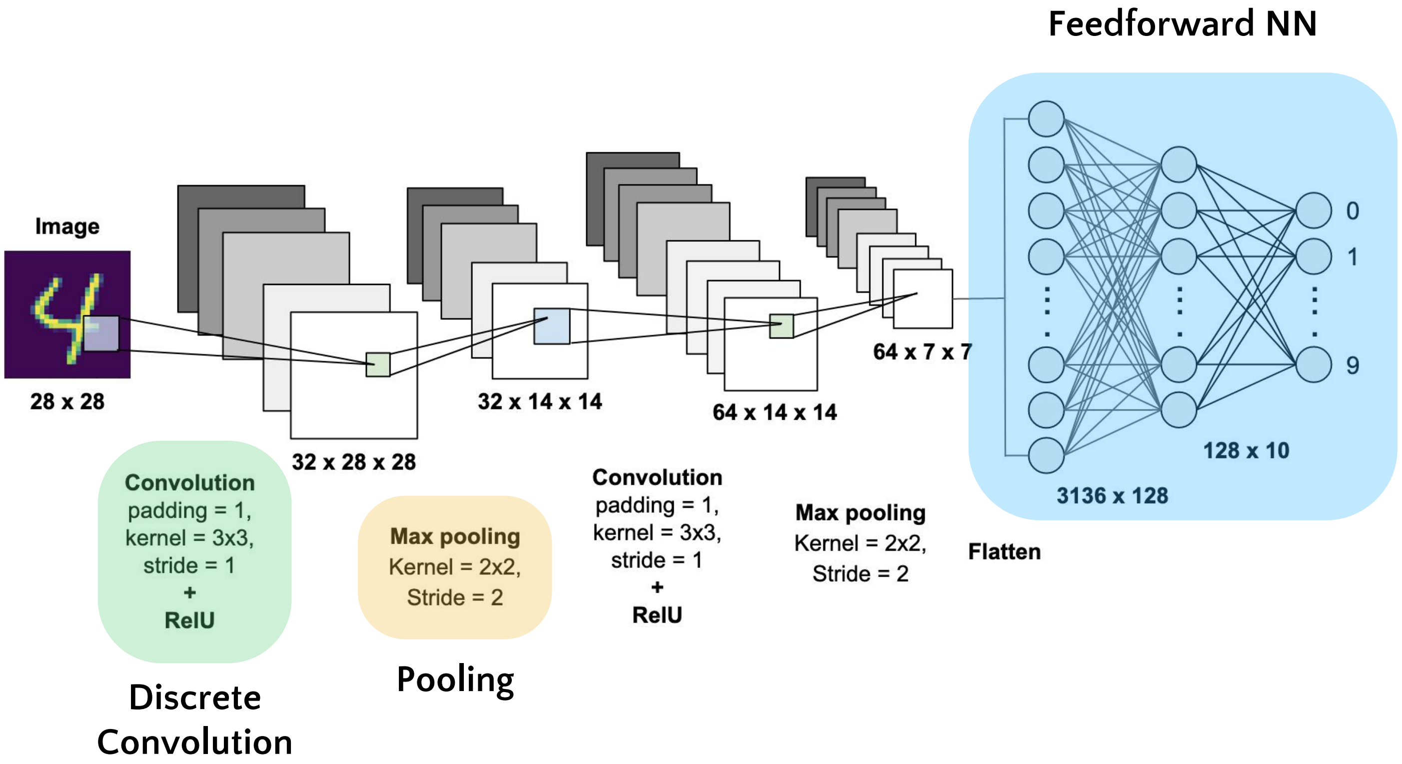

CNNs: Discrete Convolution (Correlation)

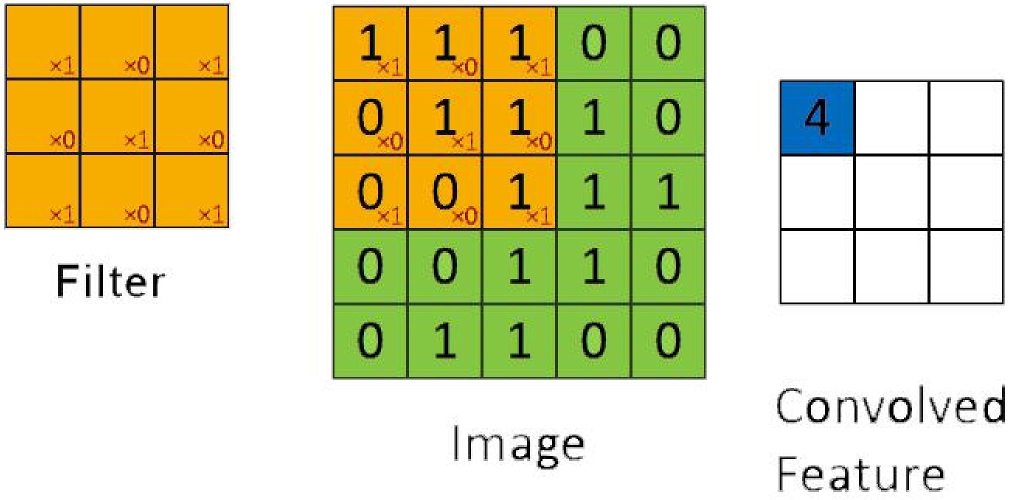

$$(f * k)(i, j) = \sum_{n=-h}^{+h} \sum_{m=-h}^{+h} f(i + n, j + m) k(i, j)$$

CNNs: (Max) Pooling

1	3	2	9
7	4	1	5
8	5	2	3
4	2	1	4

7	9
8	

$$[MaxPool(f, h, s)](i, j) = max\{f((si+1)+n, (sj+1)+m) \colon\ n, m \in \{0, ..., h-1\}\}$$

What CNNs can do

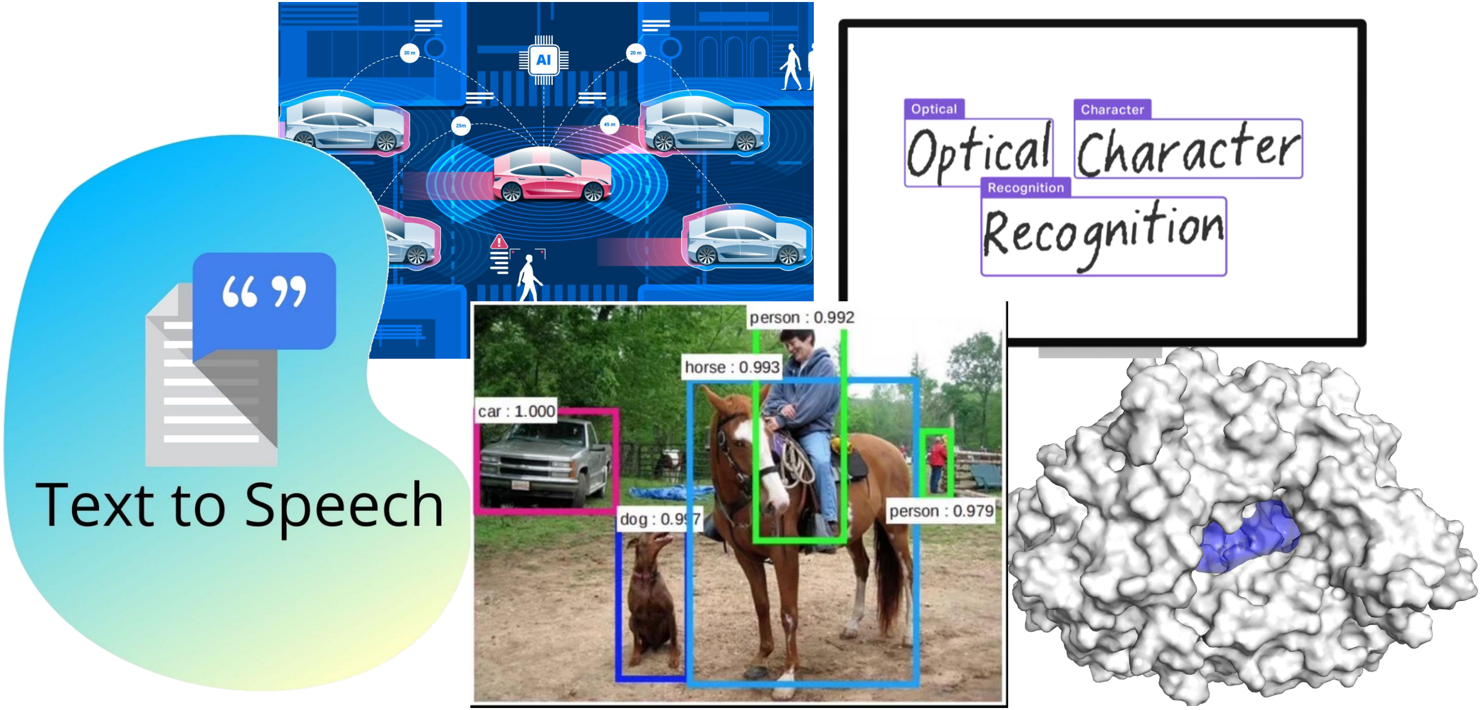

The pooling operation used in convolutional neural networks is a big mistake, and the fact that it works so well is a disaster.

Geoffrey Hinton

“

3

Deep Learning's Issues

Not all that glitters is gold

Mathematicians versus CNN Round 1

Mathematicians:

CNN: African Elephant (81.69%)

Mathematicians:

CNN: Arabian Camel (88.60%)

Mathematicians versus CNN
Round 2

Mathematicians:

CNN: African Elephant (81.69%)

Mathematicians:

CNN: LLama (31.64%)

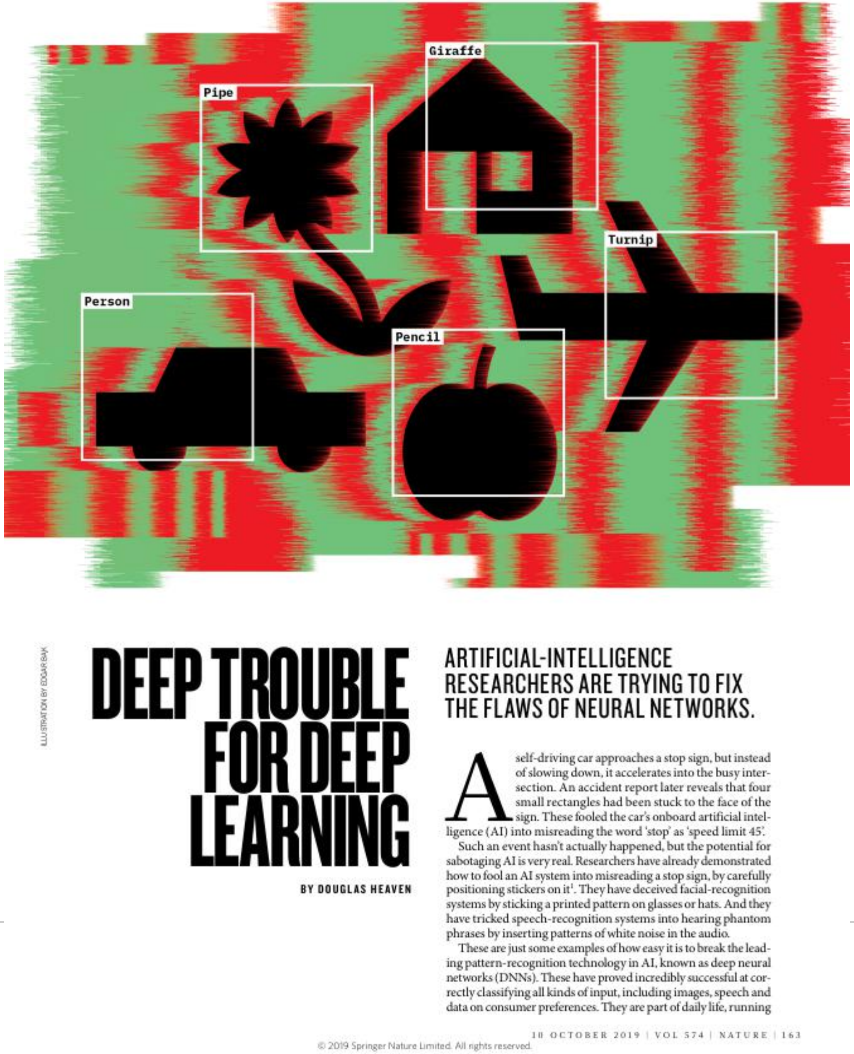

ILLUSTRATION BY EDGAR BAK

DEEP TROUBLE FOR DEEP LEARNING

ARTIFICIAL-INTELLIGENCE RESEARCHERS ARE TRYING TO FIX THE FLAWS OF NEURAL NETWORKS.

BY DOUGLAS HEAVEN

A self-driving car approaches a stop sign, but instead of slowing down, it accelerates into the busy intersection. An accident report later reveals that four small rectangles had been stuck to the face of the sign. These fooled the car's onboard artificial intelligence (AI) into misreading the word 'stop' as 'speed limit 45'.

Such an event hasn't actually happened, but the potential for sabotaging AI is very real. Researchers have already demonstrated how to fool an AI system into misreading a stop sign, by carefully positioning stickers on it[1]. They have deceived facial-recognition systems by sticking a printed pattern on glasses or hats. And they have tricked speech-recognition systems into hearing phantom phrases by inserting patterns of white noise in the audio.

These are just some examples of how easy it is to break the leading pattern-recognition technology in AI, known as deep neural networks (DNNs). These have proved incredibly successful at correctly classifying all kinds of input, including images, speech and data on consumer preferences. They are part of daily life, running

CNNs (like other architectures) have problems:

- **Unstable**
- **Non Equivariant**
- **Hackable**

Computer Scientist’s solution to second problem

(X_i, Y_i)

Data Augmentation

(X_i^1, Y_i) (X_i^2, Y_i)

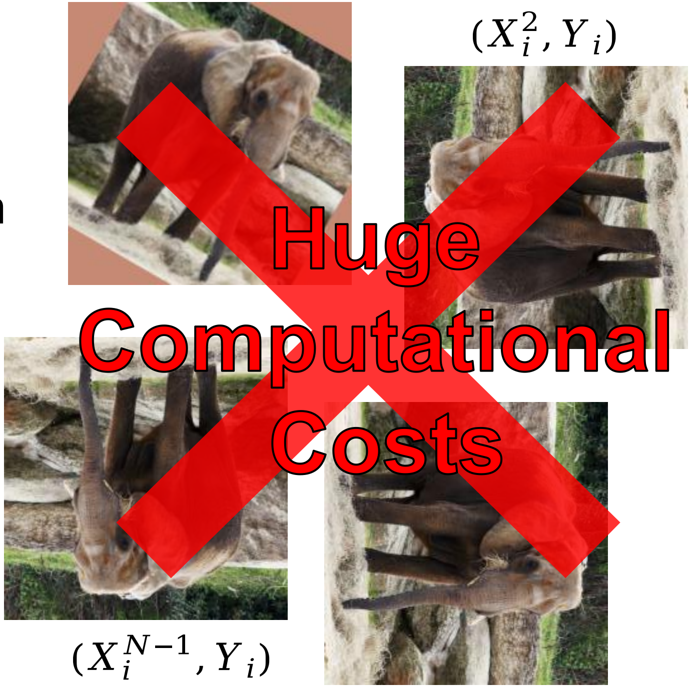

(X_i^{N-1}, Y_i) (X_i^N, Y_i)

Convolution Pros and Cons

Good

Convolution is by definition translation equivariant:

$$(f \circ \tau) * k = (f * k) \circ \tau$$

(Usually) Bad

Translations is the largest group of transformations for which this happens:

$$(f \circ \rho) * k \neq (f * k) \circ \rho$$

4

GENEOs

A possible solution

GENEOs definition

Consider two spaces of functions defined on two topological spaces X and Y

$$\Phi = \{\varphi: X \to \mathbb{R}\}$$
$$\Psi = \{\psi: Y \to \mathbb{R}\}$$

Then consider the group $Hom_{\Phi}(X)$ of all the homeomophisms of X that preserve Φ: those $g \in Hom(X)$ such that $\varphi \circ g \in \Phi$ for all $\varphi \in \Phi$.

GENEOs definition

Then chose a subgroup G of $Hom_{\Phi}(X)$. Repeat for H subgroup of $Hom_{\Psi}(X)$. Finally fix a homomorphism $T: G \to H$.

Definition (GENEOs)

A Group Equivariant Non-Expansive Operator F between (Φ, G) and (Ψ, H) wrt to T is a map $F: \Phi \to \Psi$ with the following two properties:

- **Equivariance:** $F(\varphi \circ g) = F(\varphi) \circ T(g)\ \forall \varphi, \forall g$
- **Non-Expansivity**: $\|F(\varphi_1) - F(\varphi_2)\|_\infty \leq \|\varphi_1 - \varphi_2\|_\infty$

Example

$$X = Y = \mathbb{R}^2$$
$$\Phi = \Psi = L^\infty(\mathbb{R}^2)$$
$$G = H = \{\tau_y(x) = x - y\}$$
$$T = Id_G$$
$$k \in \{h \in L^1(\mathbb{R}^2) : \|h\|_{L^1} = 1\}$$

Then the convolution operator $F(\varphi) = \varphi * k$ is a GENEO equivariant w.r.t. to planar translations.

Perspective

Exploit these (and others properties of GENEOs) to define a different kind of “Neural” Networks with broader properties of Equivariace and more robust to noise.

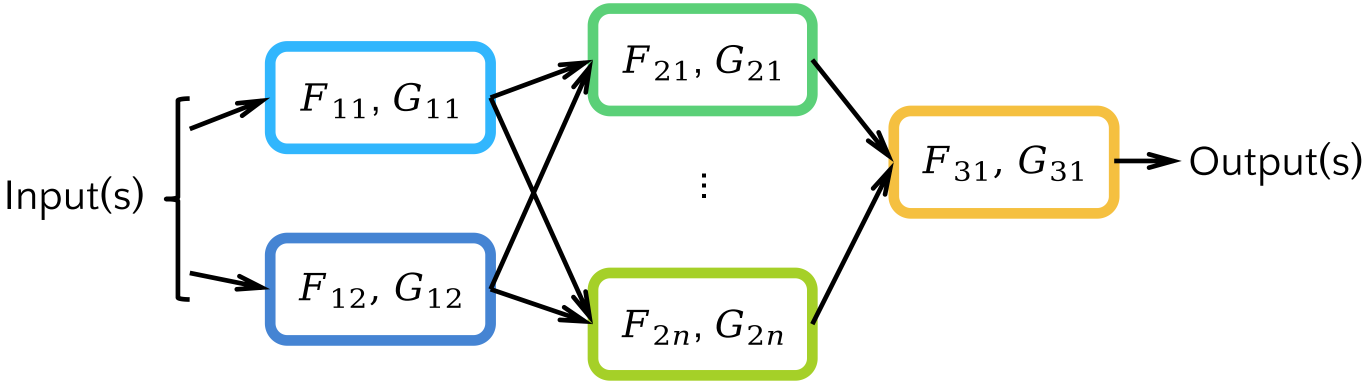

Takehome message

Neural Networks and DeepLearning are the state of the art in Artificial Intelligence. However, they are not perfect and far from beeing fully understood. Moreover mathematics can help mitigate some of the problems.

Thanks!

Any questions ?

You can reach me at

- giovanni.bocchi1@unimi.it

References
Issues with DeepLearning

- D. Heaven, "*Why deep-learning AIs are so easy to fool,*" Nature, vol. 574, no. 7777, pp. 163–166, Oct. 2019, doi: 10.1038/d41586-019-03013-5.
- C. Rudin, "*Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,*" Nature Machine Intelligence, vol. 1, no. 5, pp. 206–215, May 2019, doi: 10.1038/s42256-019-0048-x.
- S. G. Finlayson, J. D. Bowers, J. Ito, J. L. Zittrain, A. L. Beam, and I. S. Kohane, "*Adversarial attacks on medical machine learning,*" Science, vol. 363, no. 6433, pp. 1287–1289, 2019, doi: 10.1126/science.aaw4399.

References
GENEOs

- M. G. Bergomi, P. Frosini, D. Giorgi, and N. Quercioli, "*Towards a topological–geometrical theory of group equivariant non-expansive operators for data analysis and machine learning,*" Nature Machine Intelligence, vol. 1, no. 9, pp. 423–433, Sep. 2019, doi: 10.1038/s42256-019-0087-3.
- P. Frosini and N. Quercioli, "*Some remarks on the algebraic properties of group invariant operators in persistent homology,*" in 1st International Cross-Domain Conference for Machine Learning and Knowledge Extraction (CD-MAKE), Reggio, Italy, Aug. 2017, vol. LNCS-10410, pp. 14–24. doi: 10.1007/978-3-319-66808-6_2.