Geometric deep learning
- from AI to gauge theory, and back -

Daniel Persson
Department of Mathematical Sciences
Chalmers University of Technology
University of Gothenburg

Quantum gravity:
The sound of symmetry

Hermannfest

AEI Potsdam
Sept 13, 2022

Talk based on

"Geometric deep learning and equivariant neural networks"

By Gerken, Aronsson, Carlsson, Linander, Ohlsson, Petersson, **D.P.**

[arXiv: 2105.13926]

"Equivariance versus augmentation for spherical images"

By Gerken, Carlsson, Linander, Ohlsson, Petersson, **D.P.**

[arXiv: 2202.03990]

+ *work in progress*

Talk based on

"Geometric deep learning and equivariant neural networks"

By Gerken, Aronsson, Carlsson, Linander, Ohlsson, Petersson, **D.P.**

[arXiv: 2105.13926]

"Equivariance versus augmentation for spherical images"

By Gerken, Carlsson, Linander, Ohlsson, Petersson, **D.P.**

[arXiv: 2202.03990]

+ *work in progress*

Ohlsson

Carlsson

Aronsson

Petersson

Gerken

Linander

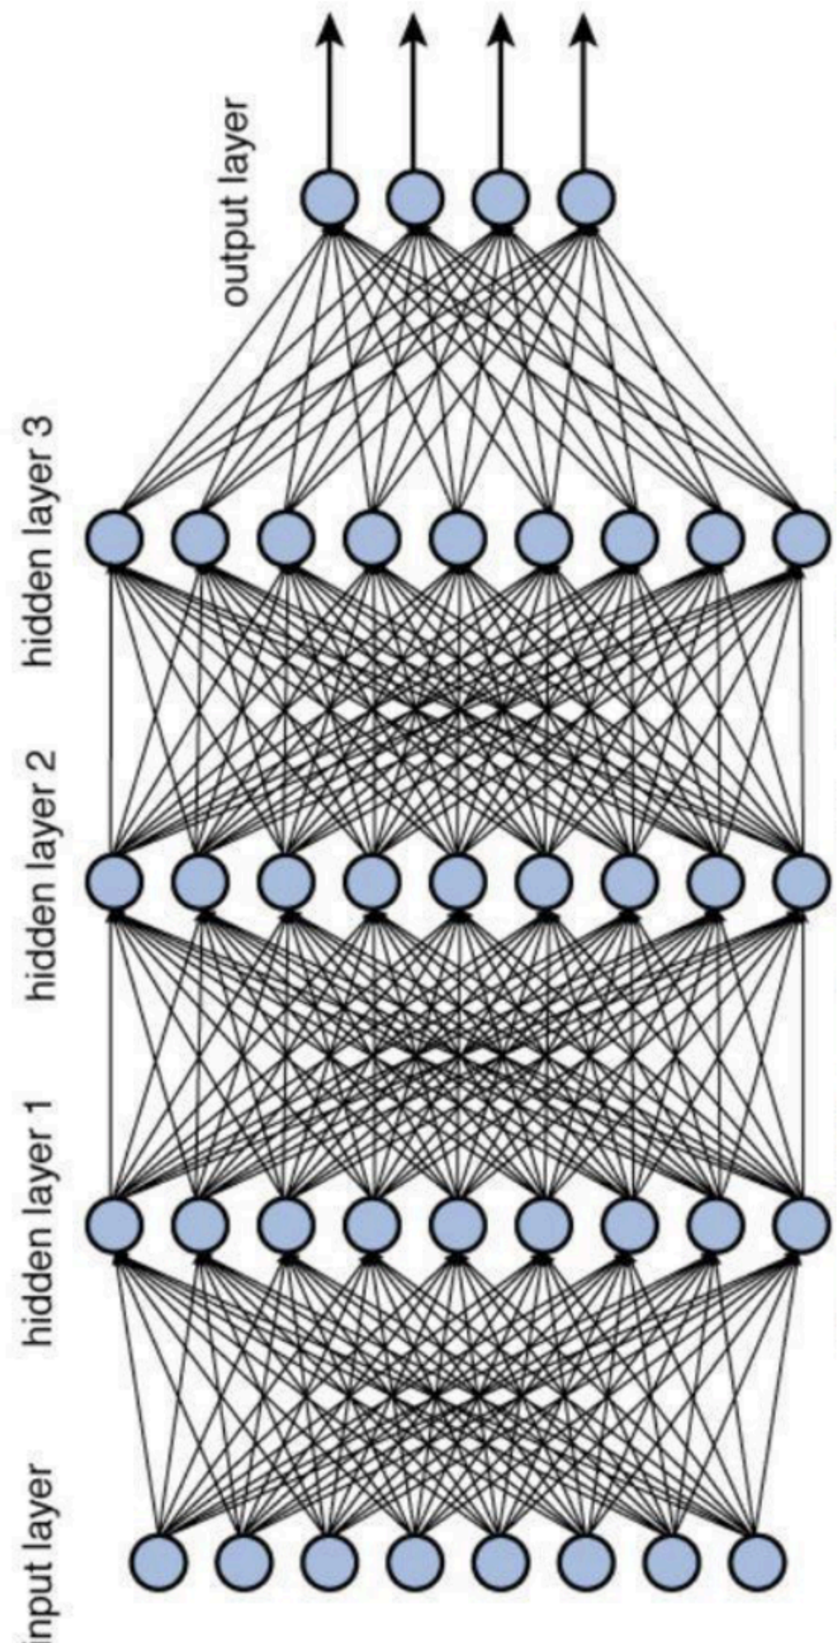

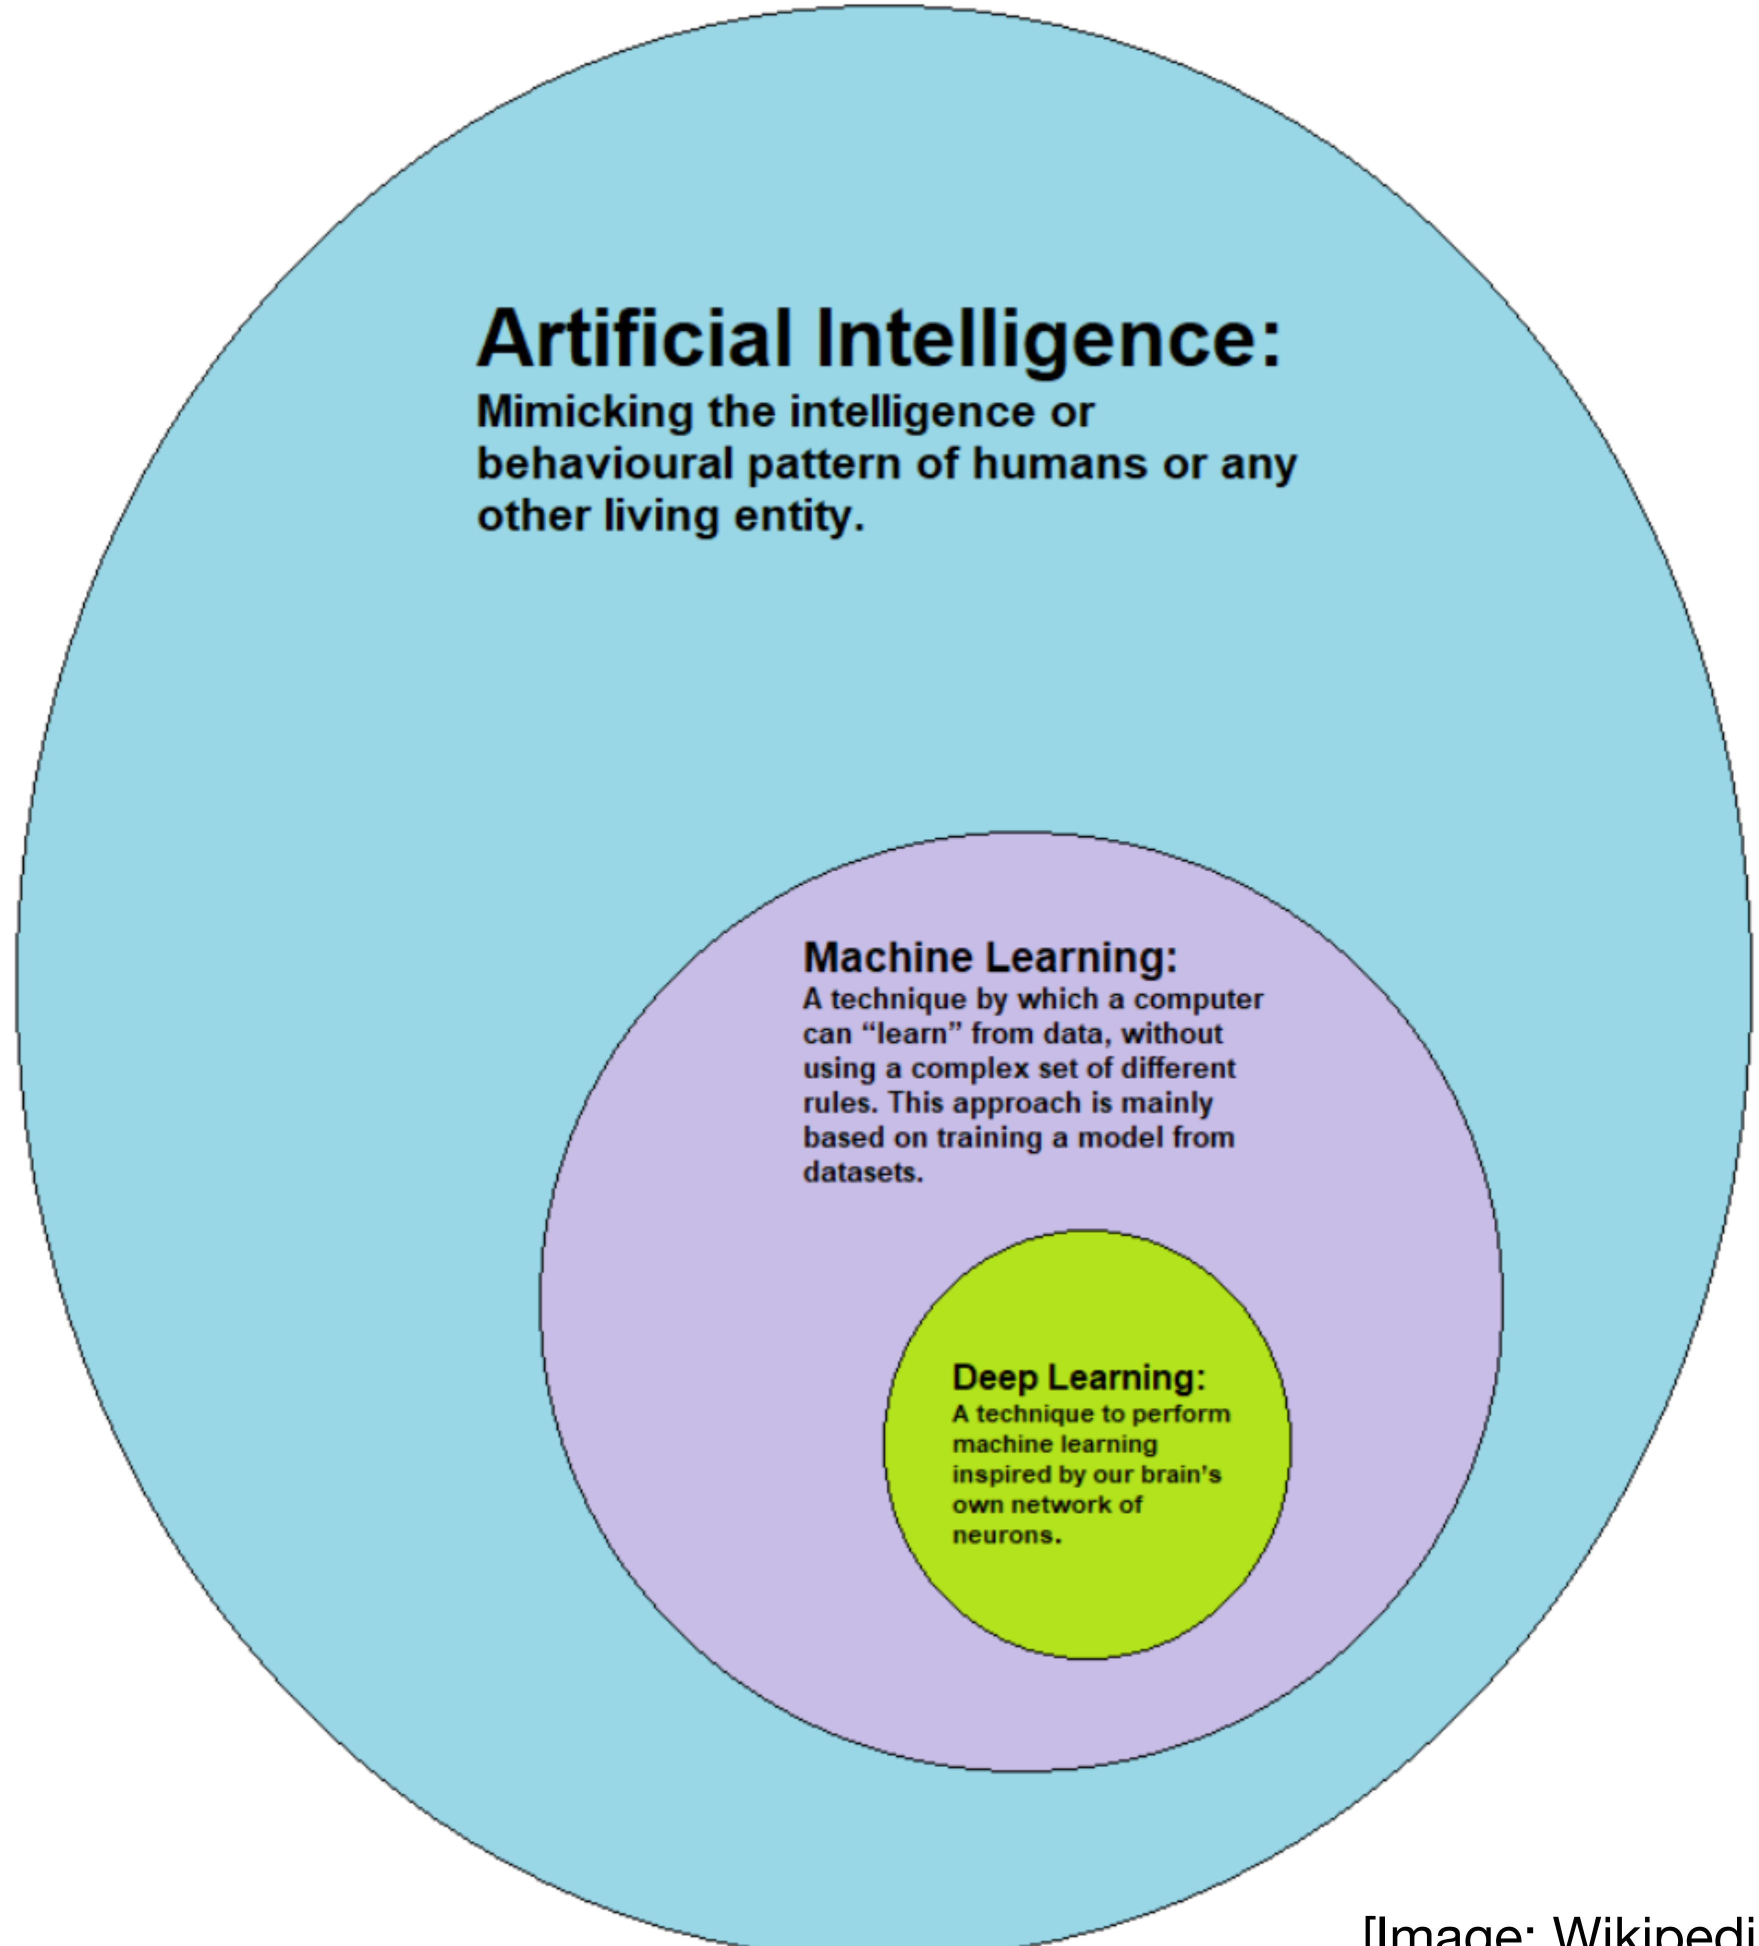

[Image: Wikipedia]

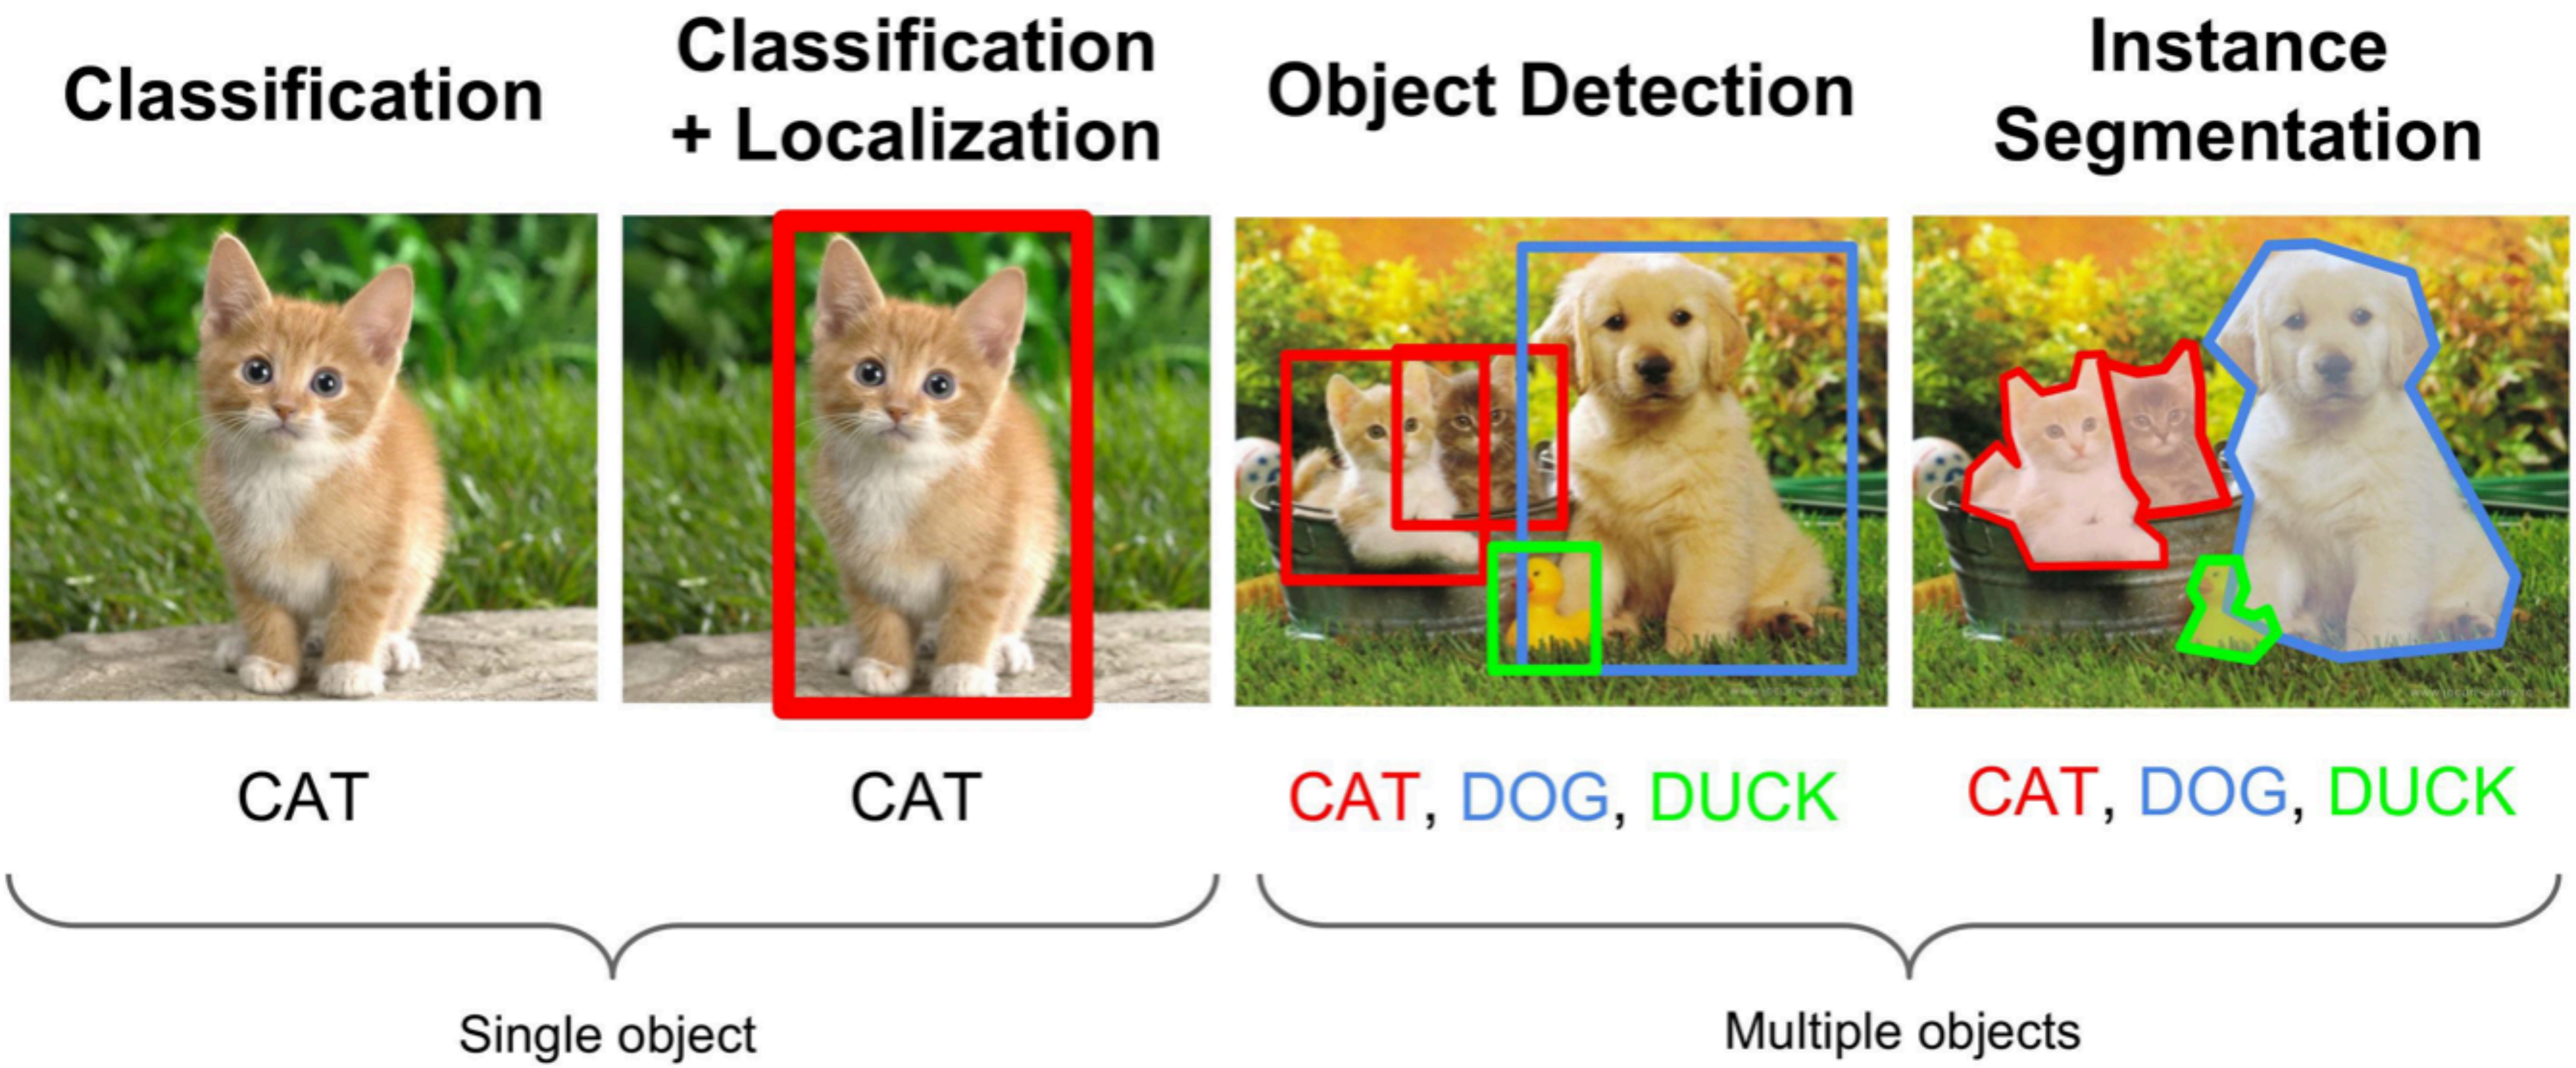

[Image from: https://medium.com/zylapp/review-of-deep-learning-algorithms-for-object-detection-c1f3d437b852]


[Bild från thispersondoesnotexist.com]


[Bild från thispersondoesnotexist.co

[Image from: https://www.walleniuswilhelmsen.com/insights/the-future-of-mobility-whats-the-road-ahead-for-self-driving-vehicles]

[Image from: https://adigaskell.org/2018/10/15/how-good-is-machine-translation/]

[Image from: https://www.walleniuswilhelmsen.com/insights/the-future-of-mobility-whats-the-road-ahead-for-self-driving-vehicles]

[Image from: https://adigaskell.org/2018/10/15/how-good-is-machine-translation/]

Artificial neuron:

1. Input layer
2. Hidden layer(s)
3. Output layer

Real Neuron

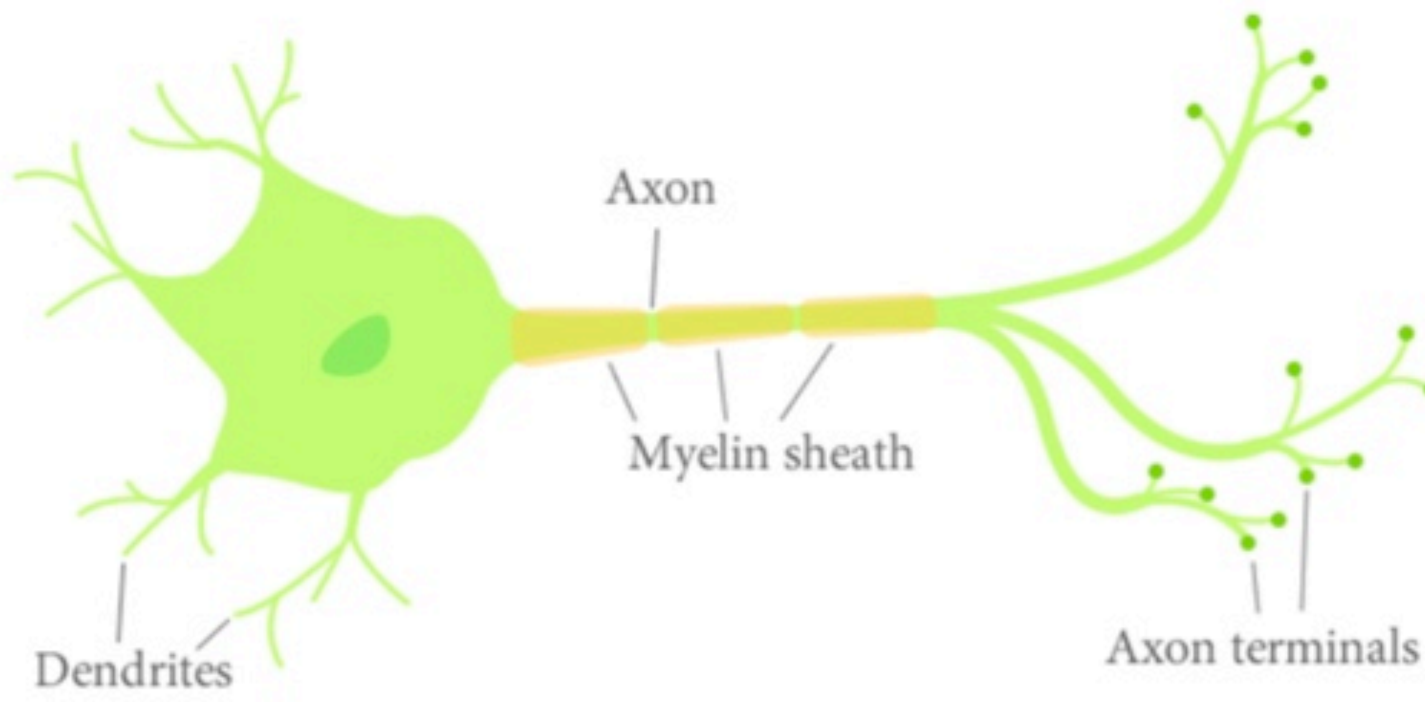

Artificial Neuron

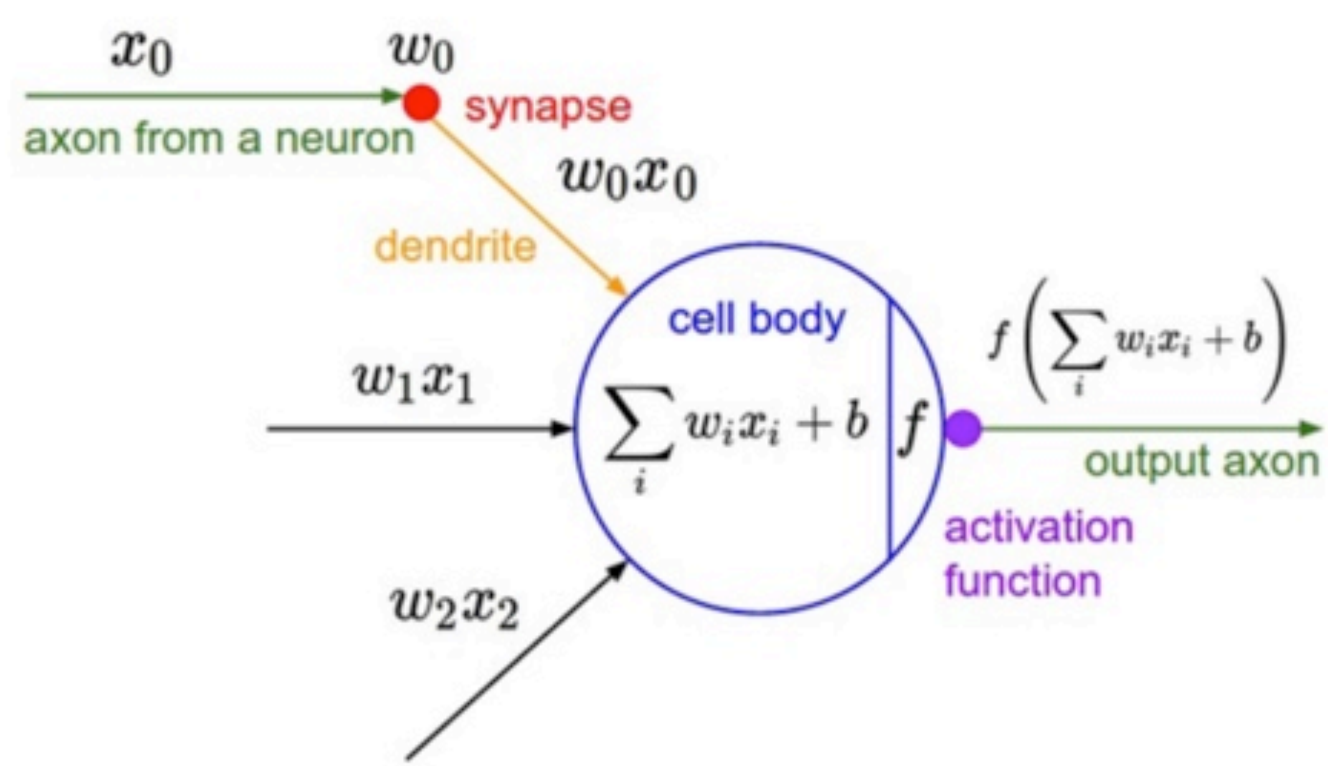

[Image from: https://towardsdatascience.com/why-is-mathematics-vital-to-thrive-in-your-ai-career-c11bd8446ddc]

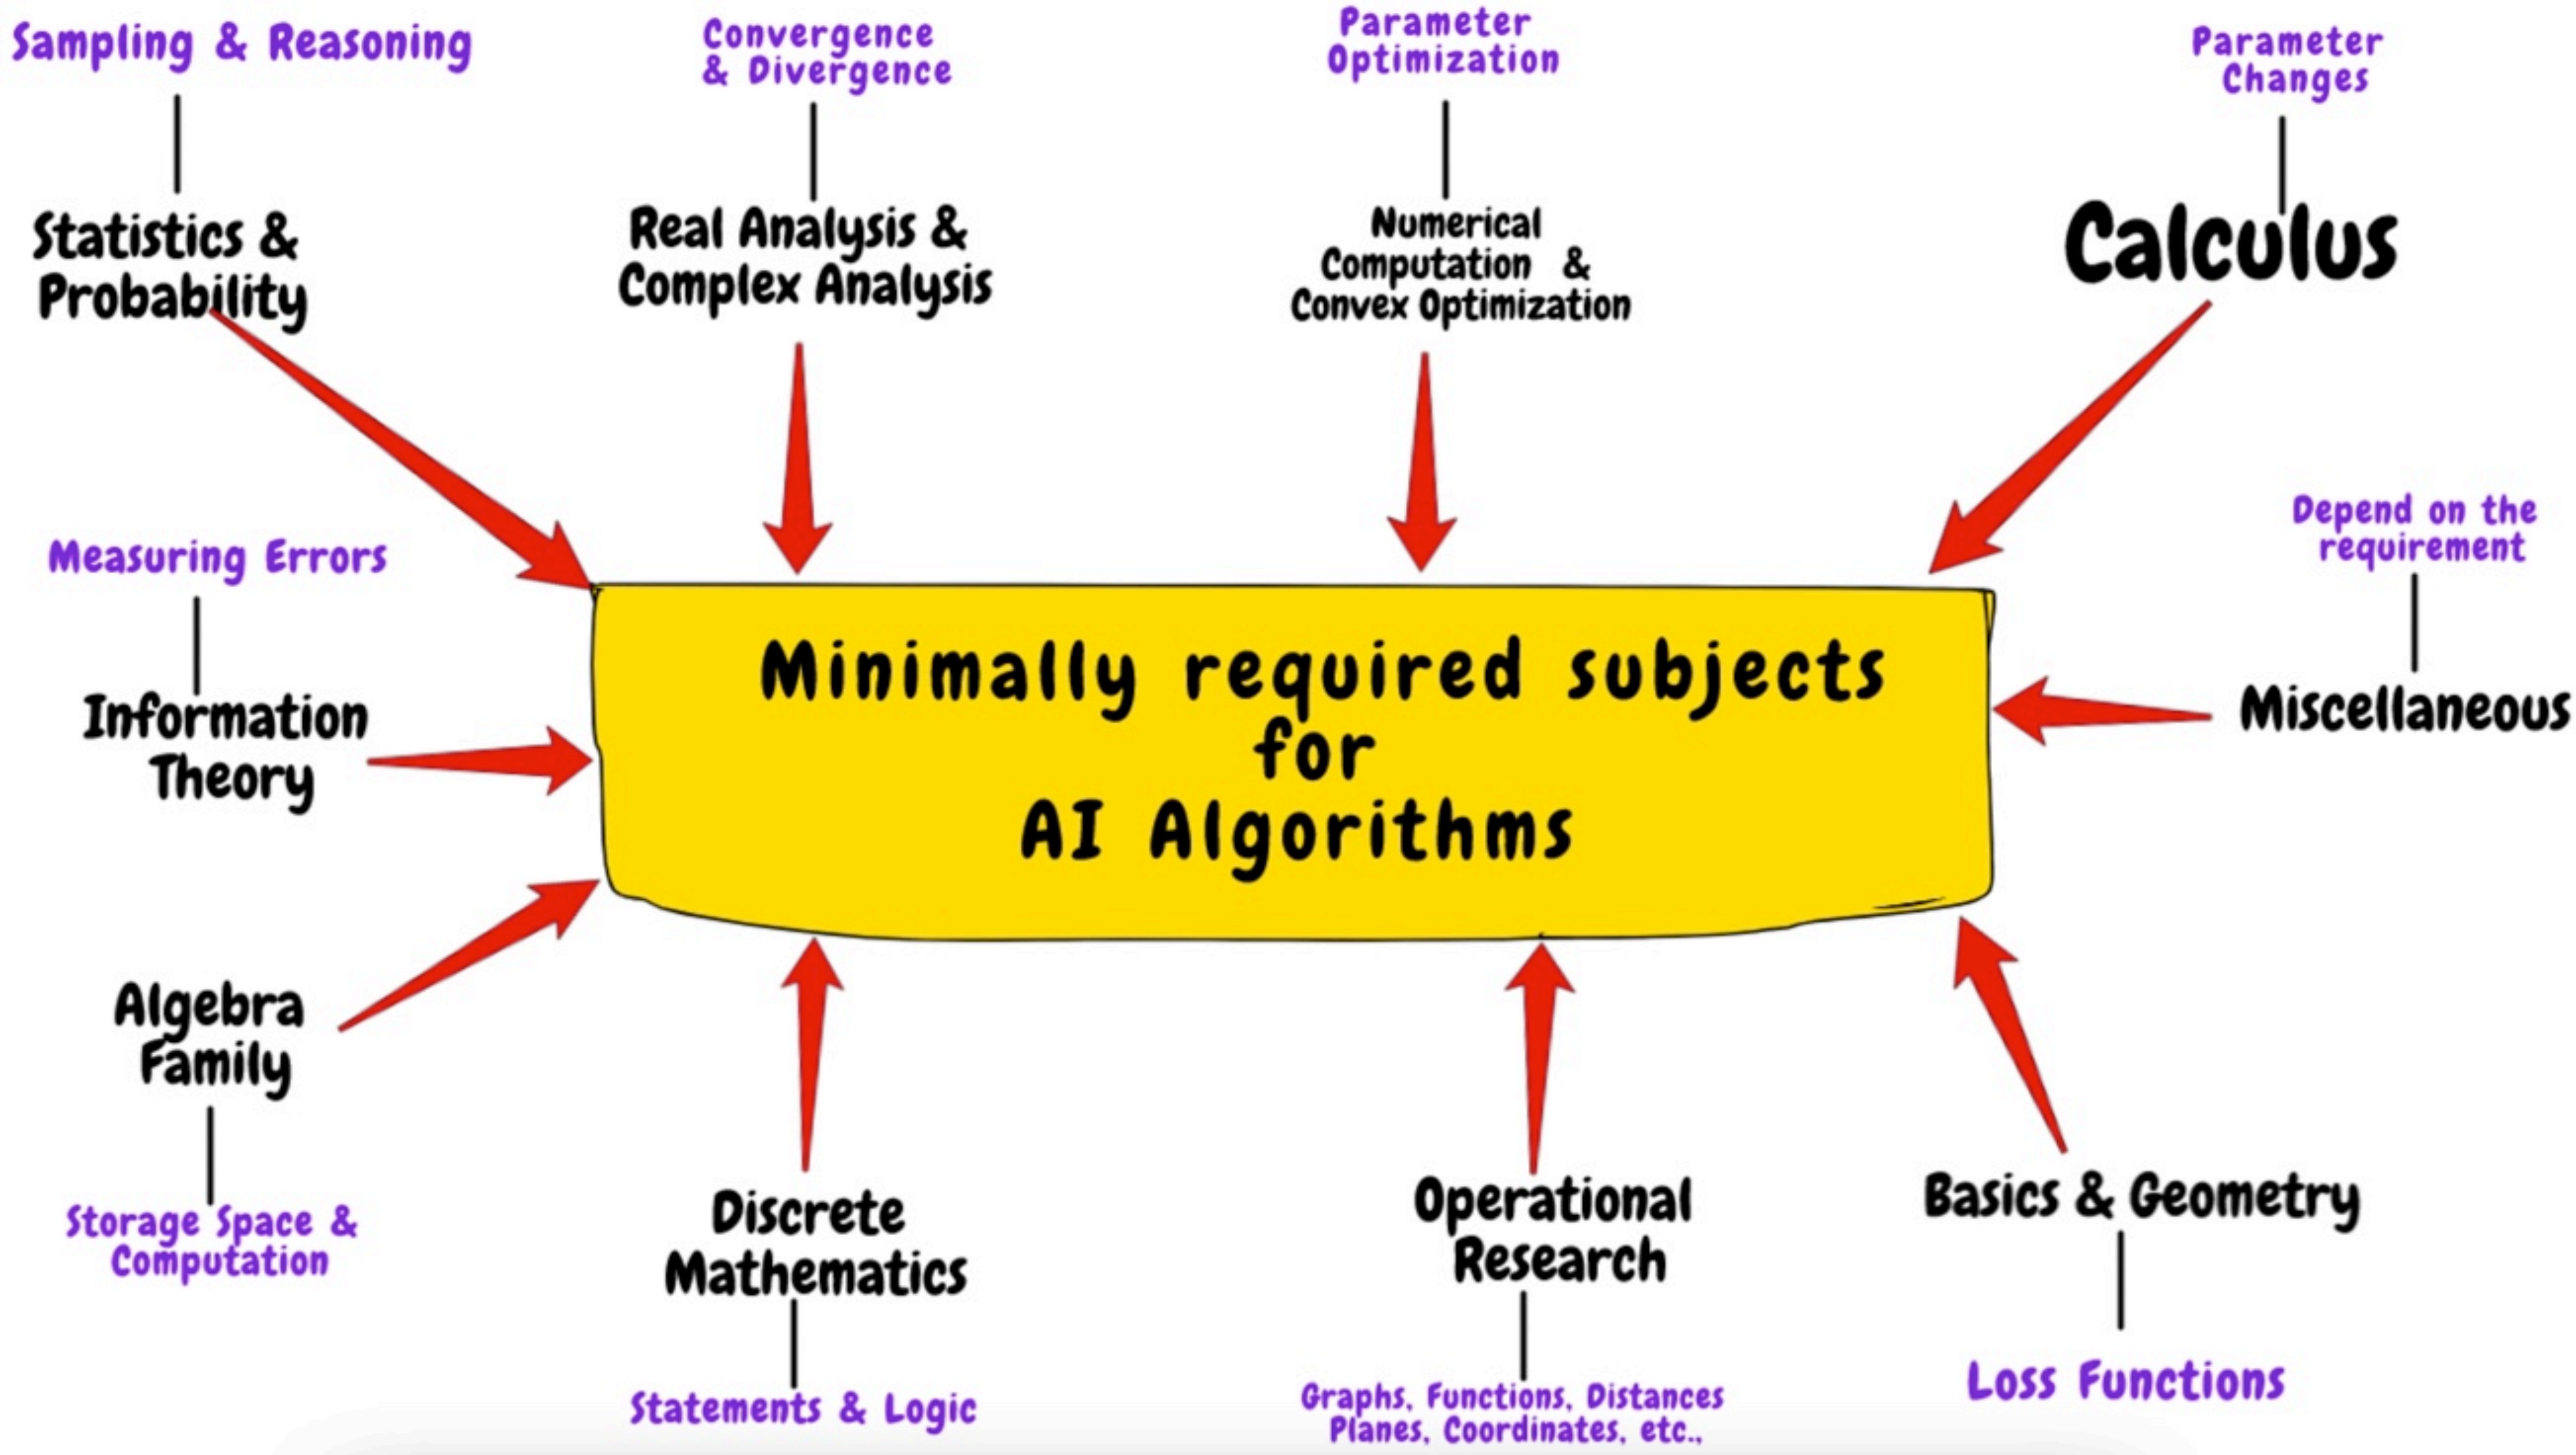
Sampling & Reasoning
Statistics & Probability
Convergence & Divergence
Real Analysis & Complex Analysis
Parameter Optimization
Numerical Computation & Convex Optimization
Parameter Changes
Calculus
Measuring Errors
Information Theory
Minimally required subjects for AI Algorithms
Depend on the requirement
Miscellaneous
Algebra Family
Storage Space & Computation
Discrete Mathematics
Statements & Logic
Operational Research
Graphs, Functions, Distances Planes, Coordinates, etc.,
Basics & Geometry
Loss Functions

Geometric Deep Learning

What to do if the data is curved?

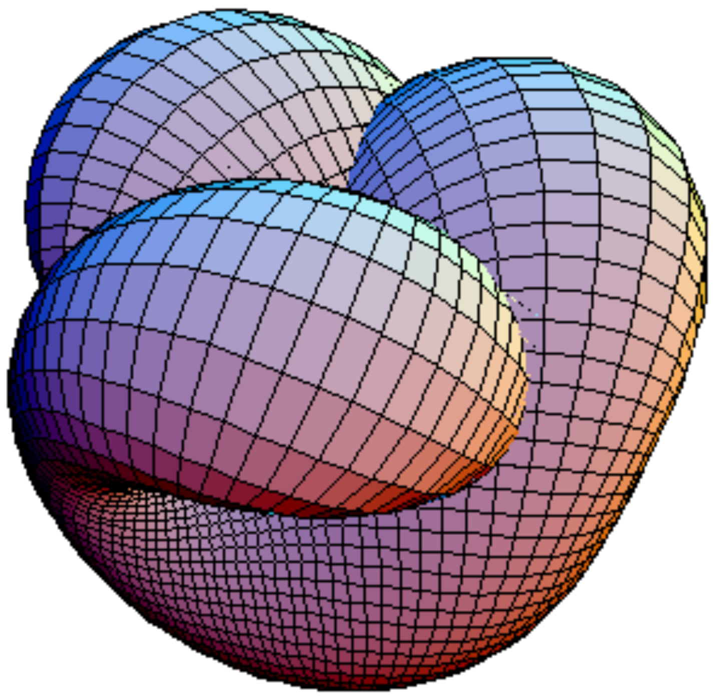

What to do if the data has symmetries?

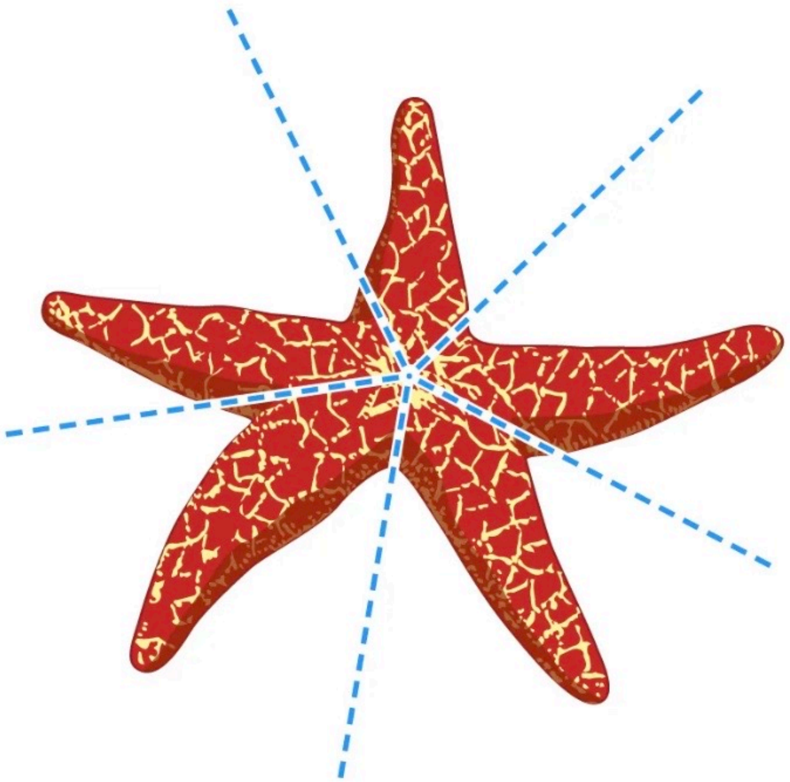

Geometric Deep Learning

Self-driving cars

Climate and weather data

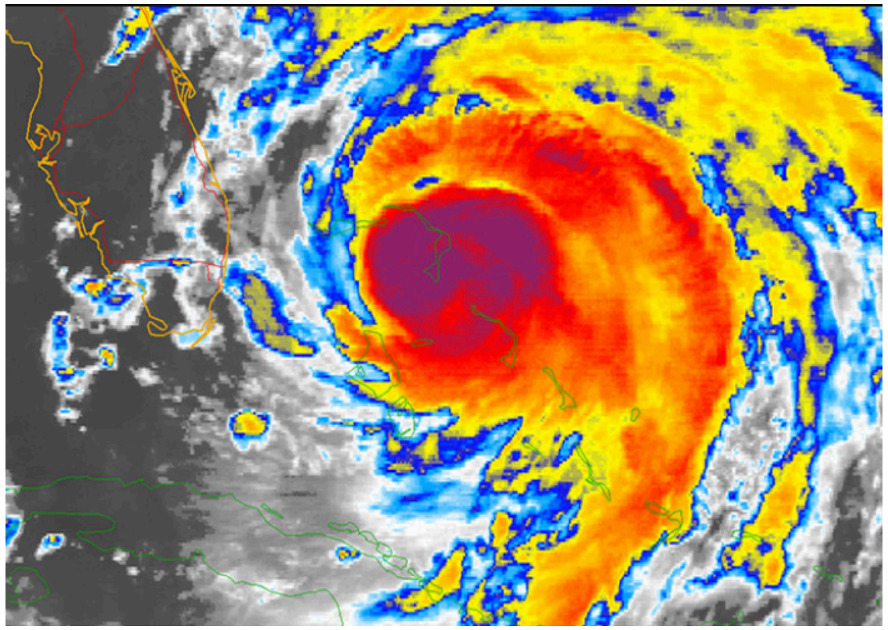

Geometric Deep Learning

But also connections to mathematics:

Deep learning on manifolds	[Cohen et al]
Sheaf neural networks	[Bronstein et al]
Graph neural networks	[LeCun et al]

And physics:

Lattice gauge theories	[Favoni et al]
Supergravity and string vacua	[He et al][Berman, Fischbacher et al]
Topological phases of matter	[in progress]

Convolutional Neural Networks

“Convolutional networks are simply neural networks that use convolution in place of general matrix multiplication in at least one of their layers.”

[Goodfellow, Bengio, Courville]

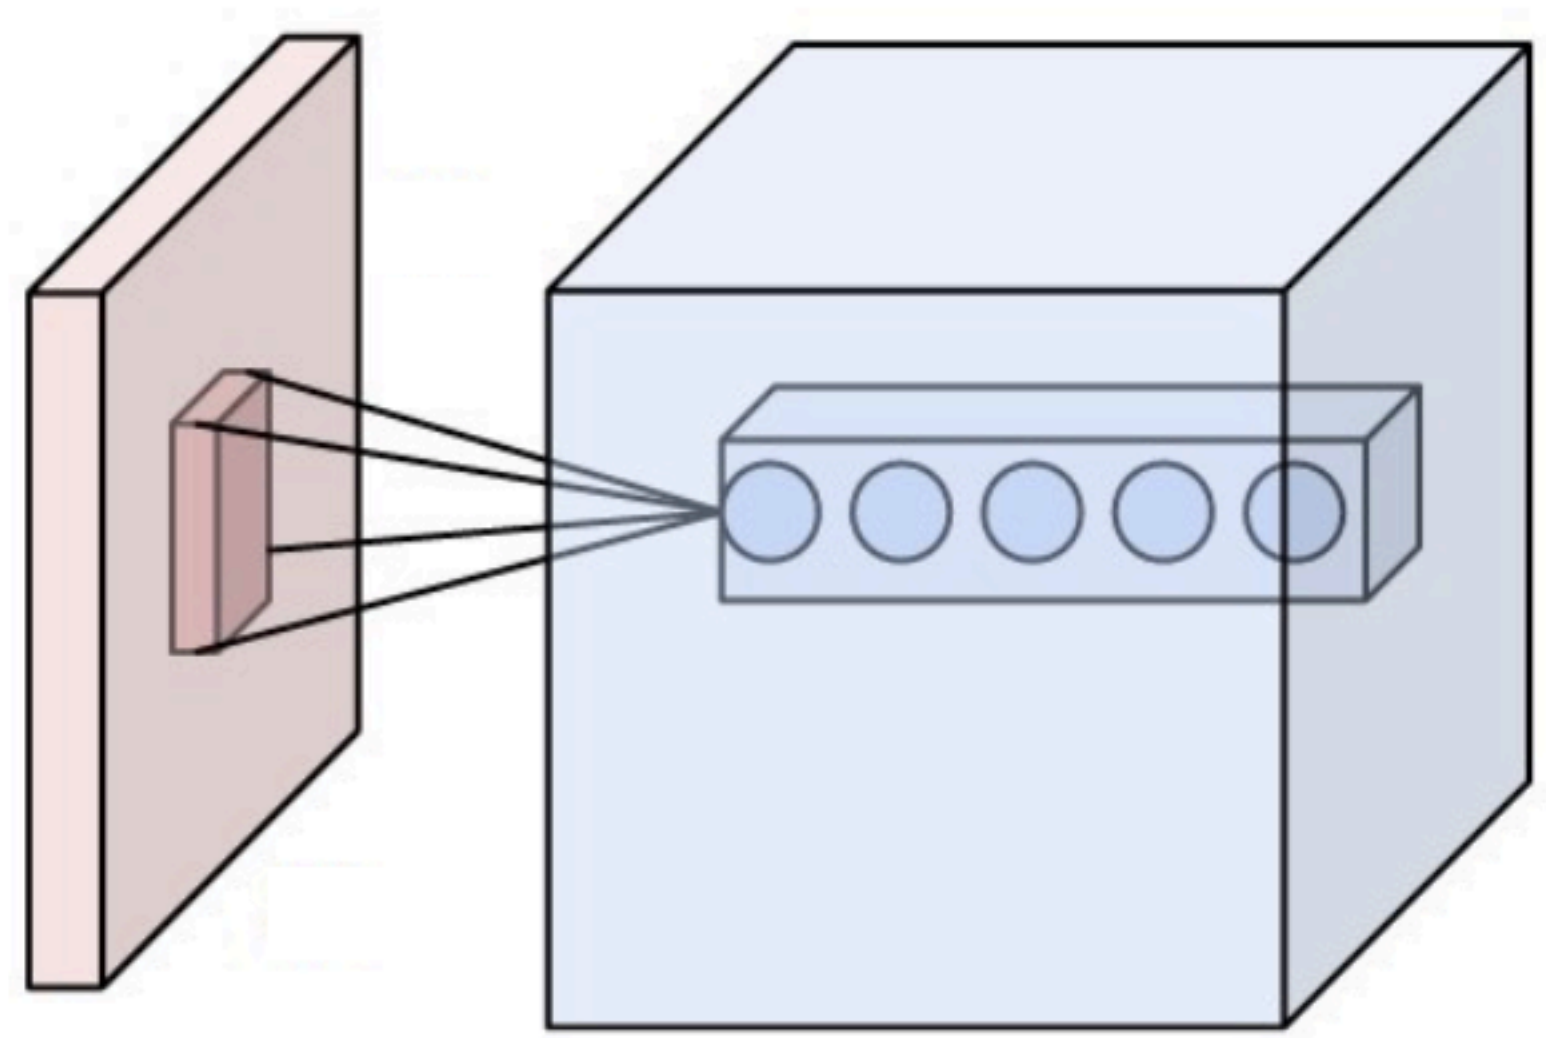

Image recognition using CNNs

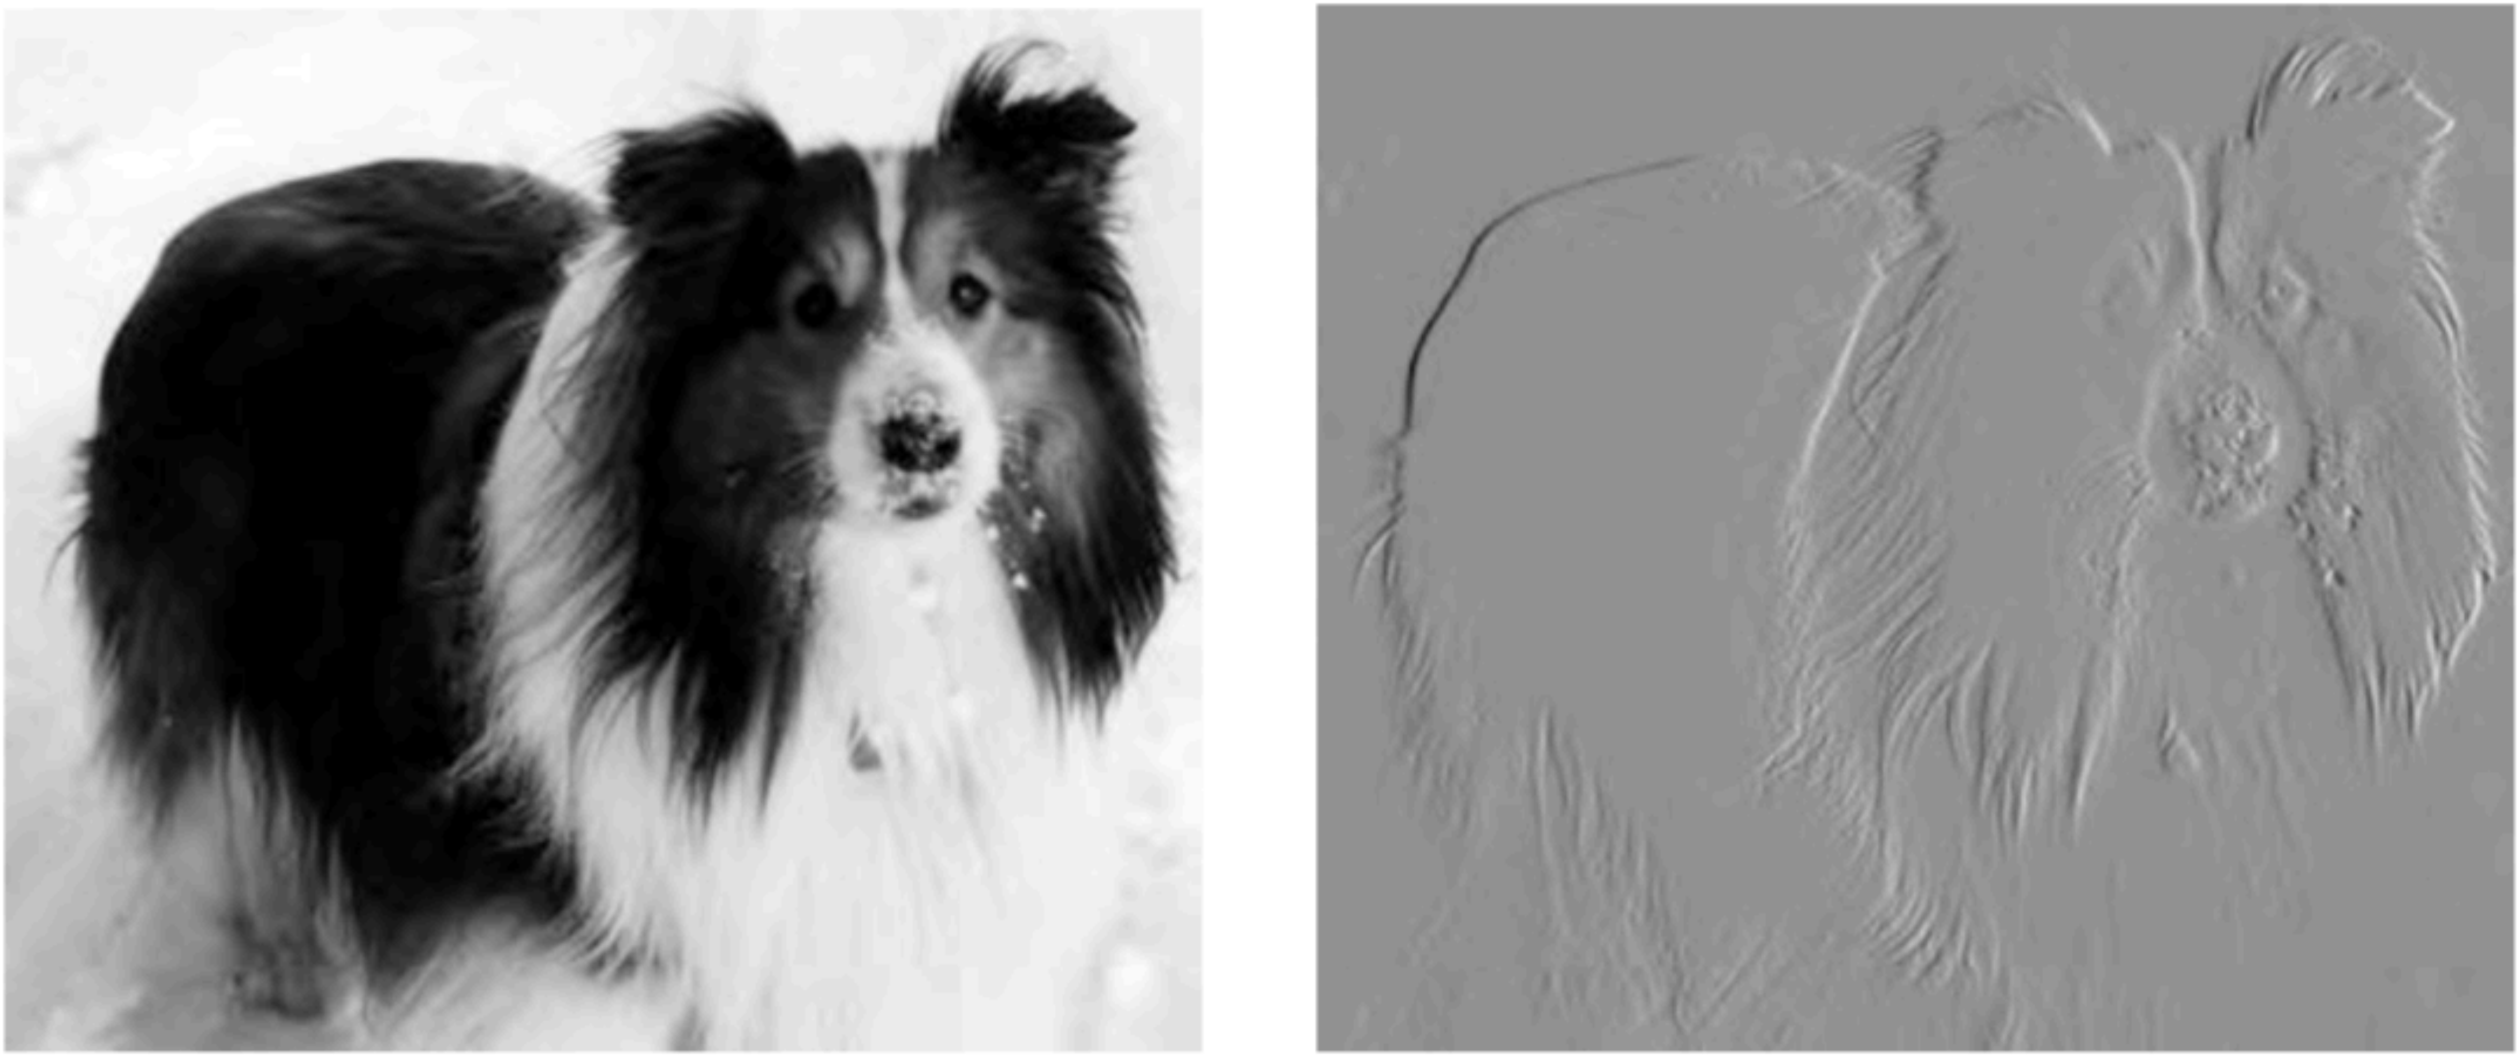

Detection using only **vertically oriented edges**. Enormous efficiency improvement compared to matrix multiplication.

[Goodfellow, Bengio, Courville]

Mathematical structure

For each layer we have a **feature map**:

$$f : \mathbb{Z}^2 \to \mathbb{R}^K$$

Mathematical structure

For each layer we have a **feature map**:

no. of channels

$$f : \mathbb{Z}^2 \to \mathbb{R}^{K}$$

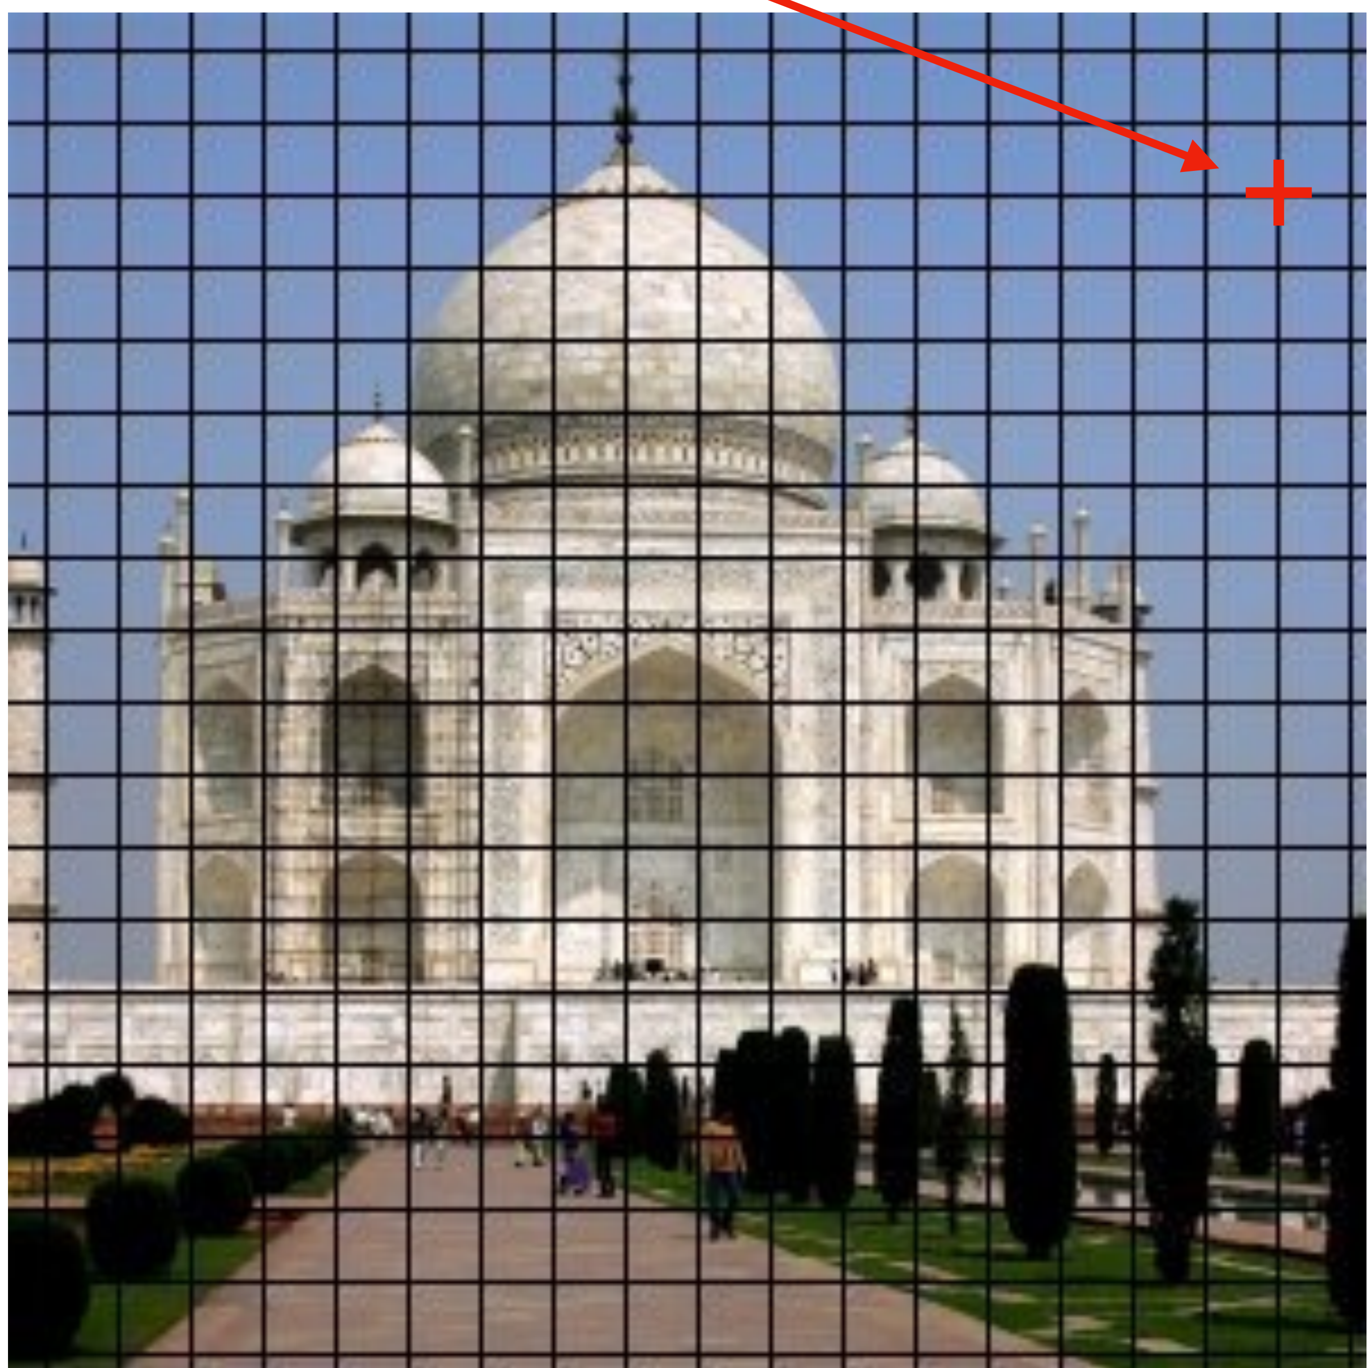

(p, q)

pixel coordinate

Kernel (filter): $\psi : \mathbb{Z}^2 \to \mathbb{R}^K$

Convolution: $$[f * \psi](x) = \sum_{y \in \mathbb{Z}^2} \sum_{k=1}^{K} f_k(y) \psi_k(x - y)$$

Kernel (filter): $\psi : \mathbb{Z}^2 \rightarrow \mathbb{R}^K$

Convolution: $$[f * \psi](x) = \sum_{y \in \mathbb{Z}^2} \sum_{k=1}^{K} f_k(y)\psi_k(x - y)$$

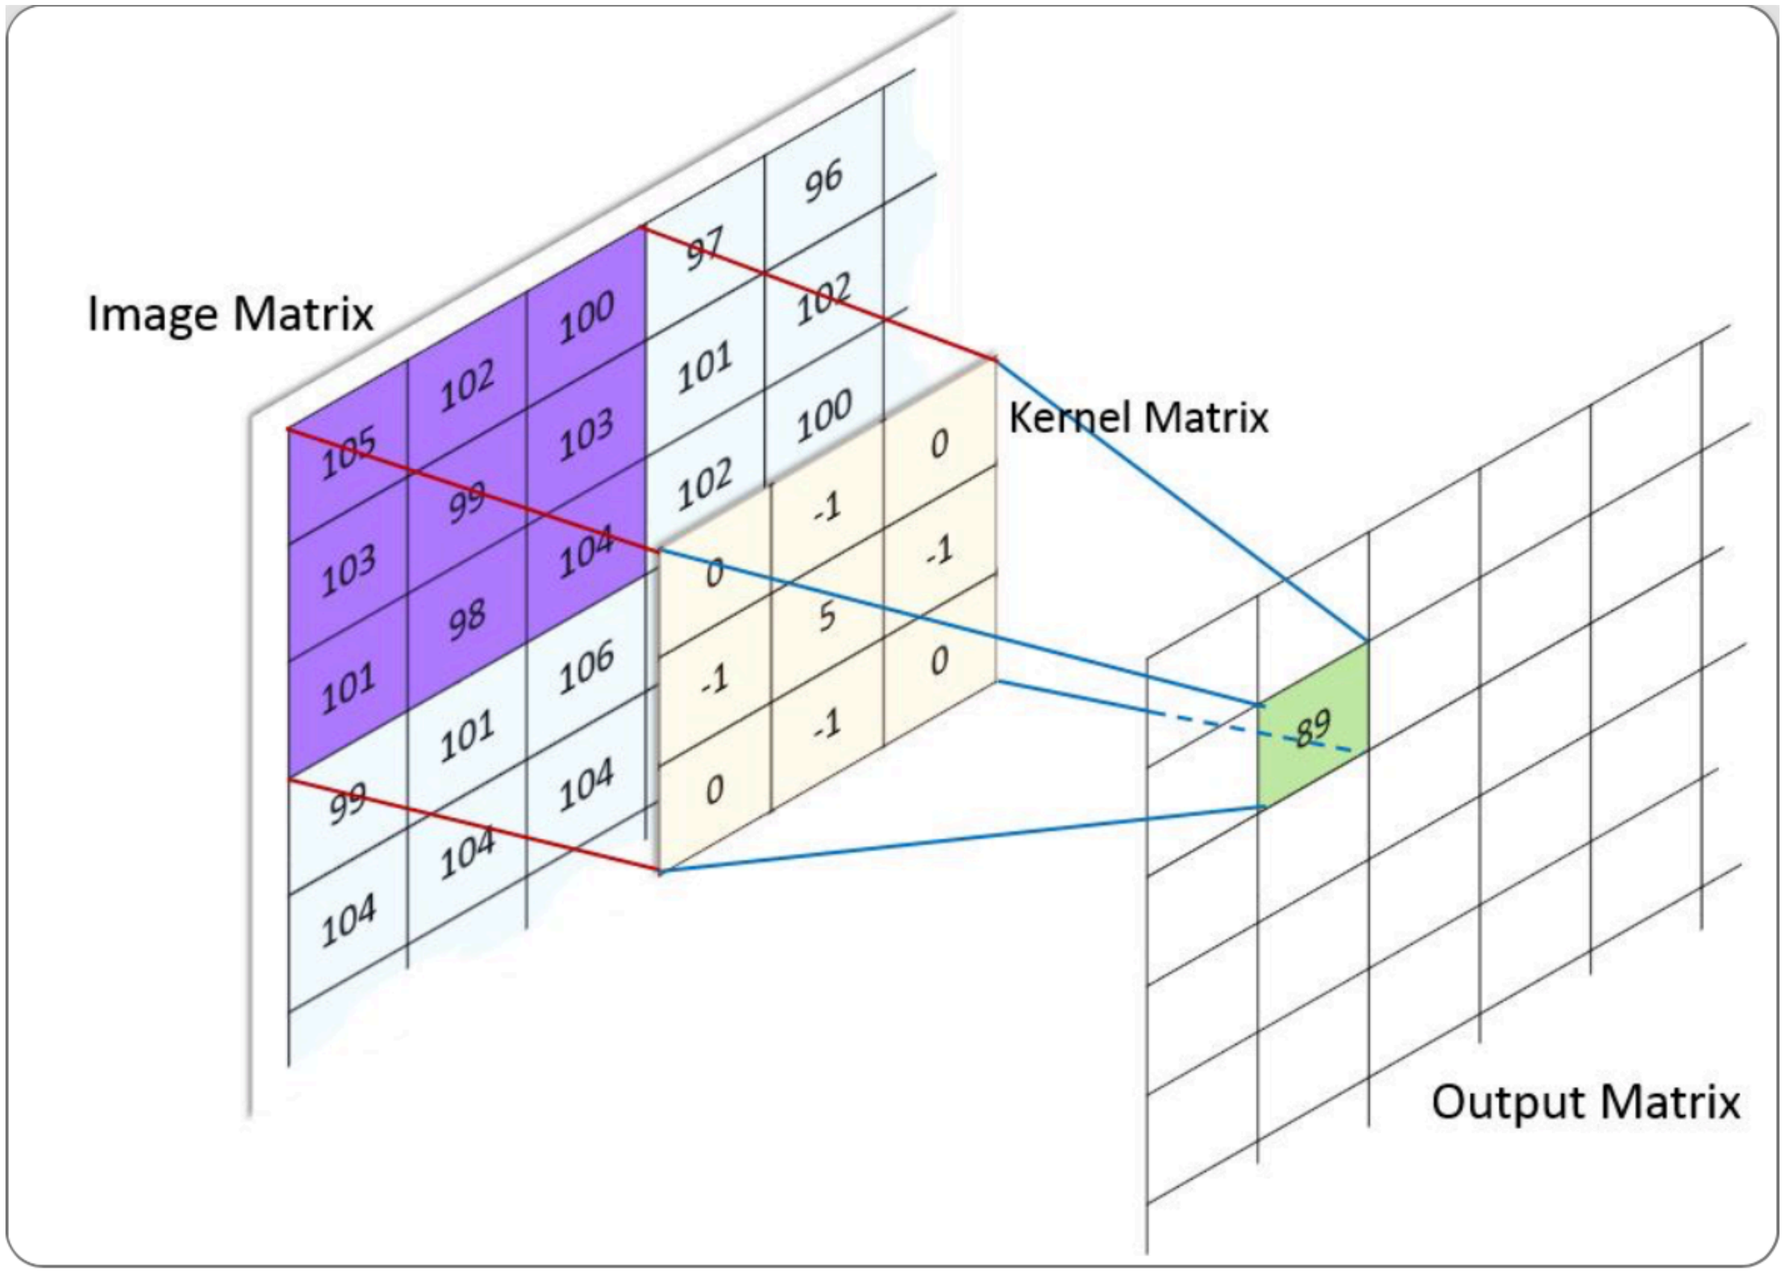

[Figure from machinelearninguru.com]

Kernel (filter): $\psi : \mathbb{Z}^2 \rightarrow \mathbb{R}^K$

Convolution: $$[f * \psi](x) = \sum_{y \in \mathbb{Z}^2} \sum_{k=1}^{K} f_k(y)\psi_k(x-y)$$

Translation map: $$[T(t)f](x) = f(x+t)$$

Kernel (filter): $\psi : \mathbb{Z}^2 \to \mathbb{R}^K$

Convolution: $$[f * \psi](x) = \sum_{y \in \mathbb{Z}^2} \sum_{k=1}^{K} f_k(y)\psi_k(x - y)$$

Translation map: $$[T(t)f](x) = f(x + t)$$

Convolution is equivariant

$$[T(t)f] * \psi = T(t)[f * \psi]$$

Convolution is equivariant

$$[T(t)f] * \psi = T(t)[f * \psi]$$

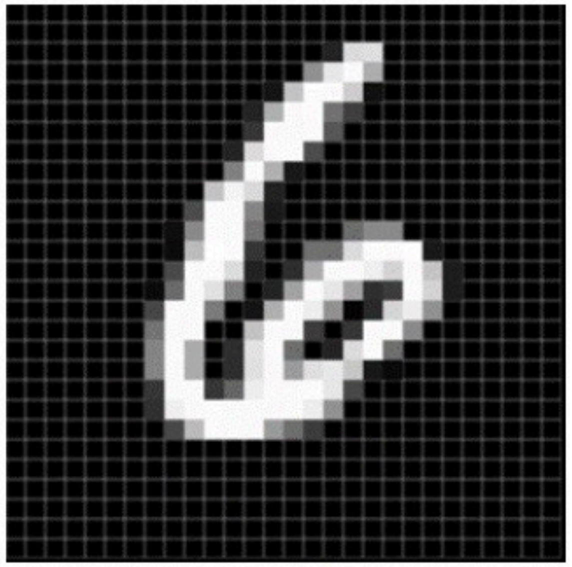

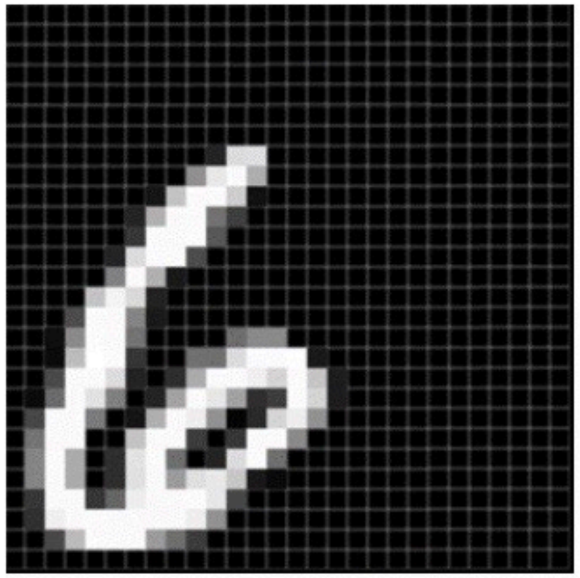

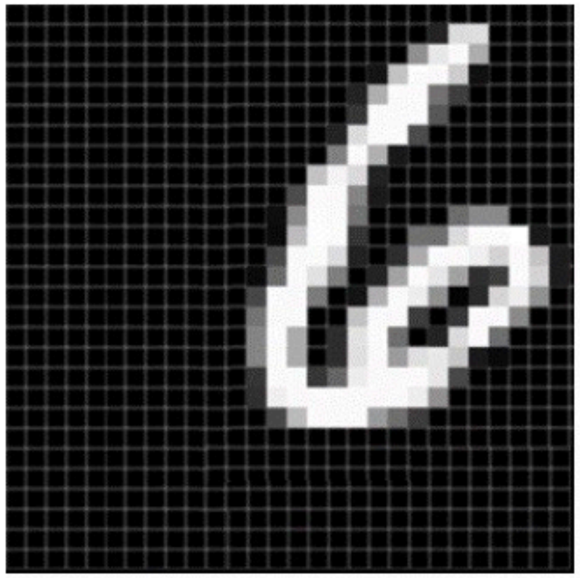

But what about more general symmetries?

Group equivariant CNNs: General framework

G a group. $H \subset G$ subgroup.

Coset space: G/H

Vector space: $V \cong \mathbb{R}^n$

Representation: $\rho : H \to GL(V)$

Group equivariant CNNs: General framework

G a group. $H \subset G$ subgroup.

Coset space: G/H

Vector space: $V \cong \mathbb{R}^n$

Representation: $\rho : H \to GL(V)$

Now consider: $P = (G \times V)/H$

This is an **equivalence class** with respect to

$$(g, v) \sim (gh, \rho(h^{-1})v)$$

This is a vector bundle:

$$
\begin{array}{ccc}
V & \longrightarrow & P \\
 & & \big\downarrow p \\
 & & G/H
\end{array}
$$

Locally, it takes the form: $G/H \times V$

Sections of P **are maps:** $s : G/H \to P$ $\quad (p \circ s = \mathrm{Id})$

Locally, we can think of these as functions $f : G/H \to V$

Locally, we can think of these as functions $f : G/H \to V$

Example: $G = \mathbb{Z}^2$ $H = \{1\}$ $V = \mathbb{R}^K$

Feature maps are sections! $f : \mathbb{Z}^2 \to \mathbb{R}^K$

Locally, we can think of these as functions $f : G/H \to V$

Example: $G = \mathbb{Z}^2$ $H = \{1\}$ $V = \mathbb{R}^K$

Feature maps are sections! $f : \mathbb{Z}^2 \to \mathbb{R}^K$

General structure of group equivariant CNNs:

$$\left\{ \begin{array}{c} \text{Feature space of a} \\ G-\text{equivariant CNN} \end{array} \right\} \cong \left\{ \begin{array}{c} \text{Sections of vector} \\ \text{bundles } P \to G/H \end{array} \right\}$$

Layers defined with group-equivariant convolutions:

$$[f * \psi](g) = \int_G \sum_{k=1}^{K} f_k(h)\psi_k(gh)dh$$

[Kondor, Trivedi][Cohen, Geiger, Weiler]

$$\left\{ \begin{array}{c} \text{Feature space of a} \\ G-\text{equivariant CNN} \end{array} \right\} \cong \left\{ \begin{array}{c} \text{Sections of vector} \\ \text{bundles } \; P \to G/H \end{array} \right\}$$

Sections of $P \to G/H$ belong to the **induced representation**:

$$\mathcal{F} = \text{Ind}_H^G \rho = \{ f : G \to V \mid f(gh) = \rho(h^{-1}) f(g) \}$$

$$\left\{ \begin{array}{c} \text{Feature space of a} \\ G-\text{equivariant CNN} \end{array} \right\} \cong \left\{ \begin{array}{c} \text{Sections of vector} \\ \text{bundles } P \to G/H \end{array} \right\}$$

Sections of $P \to G/H$ belong to the **induced representation**:

$$\mathcal{F} = \mathrm{Ind}_H^G \rho = \{ f : G \to V \mid f(gh) = \rho(h^{-1}) f(g) \}$$

$$\left\{ \begin{array}{c} \text{Maps between layers} \\ \text{in the CNN} \end{array} \right\} \cong \left\{ \begin{array}{c} G-\text{equivariant linear maps} \\ \text{between feature spaces } \mathcal{F} \to \mathcal{F}' \end{array} \right\}$$

$$\left\{ \begin{array}{c} \text{Feature space of a} \\ G-\text{equivariant CNN} \end{array} \right\} \cong \left\{ \begin{array}{c} \text{Sections of vector} \\ \text{bundles } P \to G/H \end{array} \right\}$$

Sections of $P \to G/H$ belong to the **induced representation**:

$$\mathcal{F} = \text{Ind}_H^G \rho = \{ f : G \to V \mid f(gh) = \rho(h^{-1}) f(g) \}$$

$$\left\{ \begin{array}{c} \text{Maps between layers} \\ \text{in the CNN} \end{array} \right\} \cong \left\{ \begin{array}{c} G-\text{equivariant linear maps} \\ \text{between feature spaces } \mathcal{F} \to \mathcal{F}' \end{array} \right\}$$

$$\text{Hom}_G(\mathcal{F}, \mathcal{F}')$$

(intertwining operators)

Example: Spherical signals

By Gerken, Carlsson, Linander, Ohlsson, Petersson, **D.P.**

[arXiv: 2202.03990]

$$G = SO(3)$$

$$H = SO(2)$$

$$G/H \cong S^2$$

Feature maps

$$f : S^2 \to \mathbb{R}^K$$

Relevant for :

- Omnidirectional vision
- Weather and climate data
- Cosmology & astrophysics

$$(\kappa \star f)(R) = \int_{S^2} \kappa(R^{-1}x) f(x)\ \mathrm{d}x$$

$$(\kappa \star f)(R) = \int_{\mathrm{SO}(3)} \kappa(S^{-1}R) f(S)\, \mathrm{d}S$$

Equivariance versus Augmentation for Spherical Images

JAN E. GERKEN[1], OSCAR CARLSSON[1], HAMPUS LINANDER[2], FREDRIK OHLSSON[3], CHRISTOFFER PETERSSON[1,4] AND DANIEL PERSSON[1]

Abstract

We analyze the role of rotational equivariance in convolutional neural networks (CNNs) applied to spherical images. We compare the performance of the group equivariant networks known as S2CNNs and standard non-equivariant CNNs trained with an increasing amount of data augmentation. The chosen architectures can be considered baseline references for the respective design paradigms. Our models are trained and evaluated on single or multiple items from the MNIST or FashionMNIST dataset projected onto the sphere. For the task of image classification, which is inherently rotationally *invariant*, we find that by considerably increasing the amount of data augmentation and the size of the networks, it is possible for the standard CNNs to reach at least the same performance as the equivariant network. In contrast, for the inherently *equivariant* task of semantic segmentation, the non-equivariant networks are consistently outperformed by the equivariant networks with significantly fewer parameters. We also analyze and compare the inference latency and training times of the different networks, enabling detailed tradeoff considerations between equivariant architectures and data augmentation for practical problems. The equivariant spherical networks used in the experiments will be made available at https://github.com/JanEGerken/sem_seg_s2cnn.

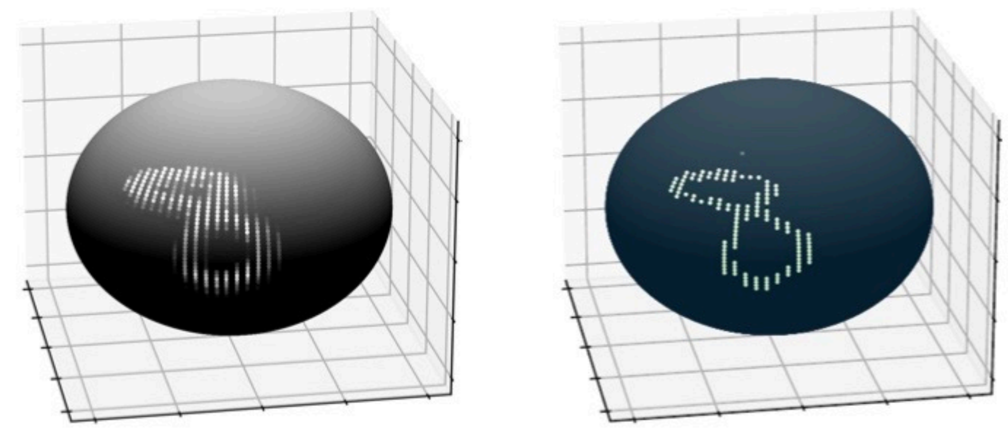

FIGURE 1.1. Sample from the spherical MNIST dataset used for semantic segmentation. Left: input data. Right: segmentation mask.

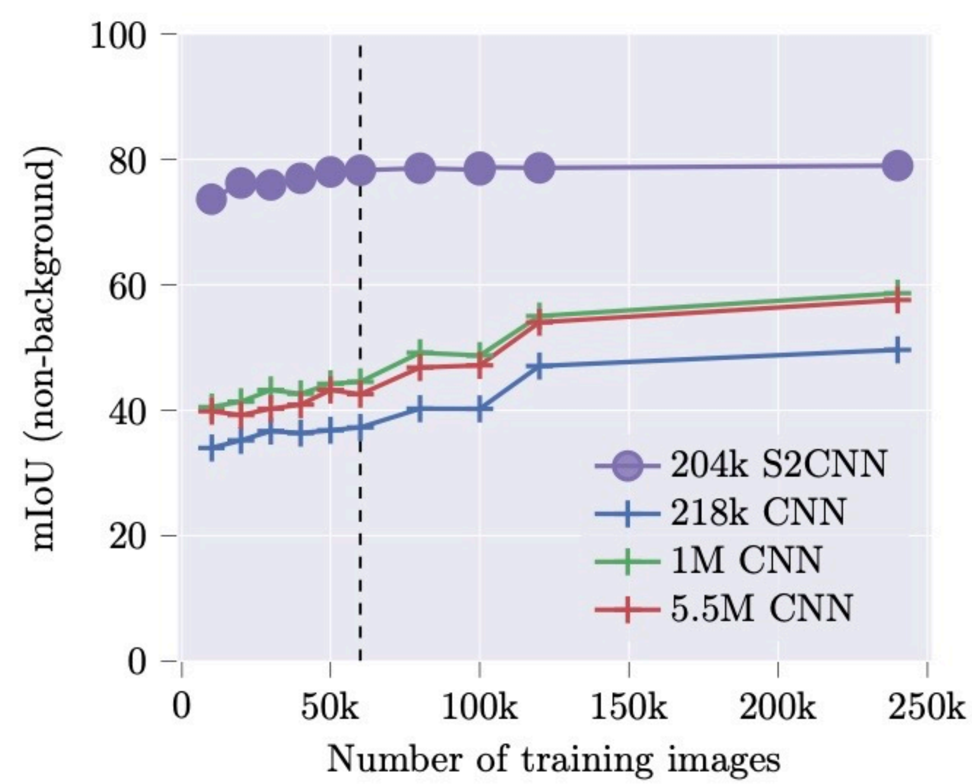

FIGURE 1.2. *Semantic Segmentation on Spherical*

Work in progress: Equivariant neural networks for autonomous driving

w/ Gerken, Carlsson, Aronsson, Linander, Ohlsson, Petersson, **D.P.**

[Image from the dataset Woodscape, projected onto the sphere]

Gauge equivariant neural networks

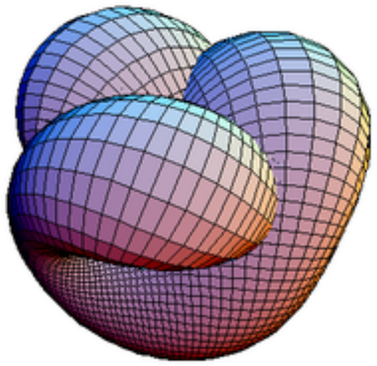

[Cheng, Anagiannis, Weiler, de Haan, Cohen, Welling]

[Gerken, Carlsson, Aronsson, Linander, Ohlsson, Petersson, **D.P.**]

CNNs on arbitrary manifolds require local equivariance	covariance w. r. t. **gauge transformations** (general coordinate transformations)
gauge equivariant feature maps	**Fields** Sections of vector bundles (frame bundles)
"elementary feature types" ?	irreducible representations of G **elementary particles** (scalars, vectors, spinors…)

Are these the seeds of a deeper relation between neural networks and gauge theory?

Outlook

The **cross-fertilization** between deep learning, theoretical physics and mathematics is an exciting rapidly developing area of research

- Is renormalization a universal principle for deep learning?
- Relation with Quantum Information Theory?
- Can we realize a neural network as a (quantum) dynamical system?
- Relation with optimal transport theory and information geometry?
- Can we implement symmetries and conservation laws?
- A spacetime perspective of Deep Neural Networks?
- Emergent phenomena?

E_{10} - equivariant neural networks?

Use deep learning to calculate root multiplicities in E_{10}?

- E_{10} - equivariant neural networks?

- Use deep learning to calculate root multiplicities in E_{10}?