Thinking Humanly	**Thinking Rationally**
"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense." (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of mental faculties through the use of computational models." (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Acting Humanly	**Acting Rationally**
"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)	"Computational Intelligence is the study of the design of intelligent agents." (Poole *et al.*, 1998) "AI . . . is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

MUNI

Rychlý průlet historií AI

PSBB084

Vojtěch Formánek

Filozofická fakulta, Masarykova Univerzita

February 15, 2024

The need

- AI as a field is a LEGO
- AI models are LEGOs
- modeling cognition and language
- planning, banking, navigation, games

Practical need

- practically incomputable (NP-hard) problems
- Traveling Salesman Problem (TSP)
- Brno-venkov has 187 villages $\implies$ $187! \sim 10^{345}$ possibilities
- number of atoms in the universe $\sim 10^{100}$

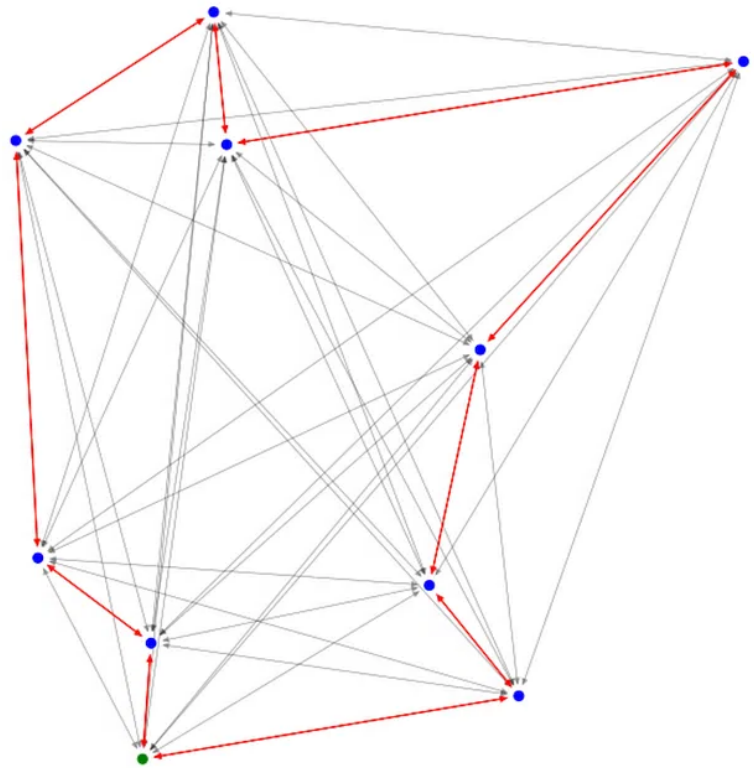

Practical need

1000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000
0000000000000000000000000000000000000000000000000
0000000000000000000

State spaces

- car travels to three nearest villages
- $\sim 3^{187}$ solutions
- state space for three visited villages
- a search tree

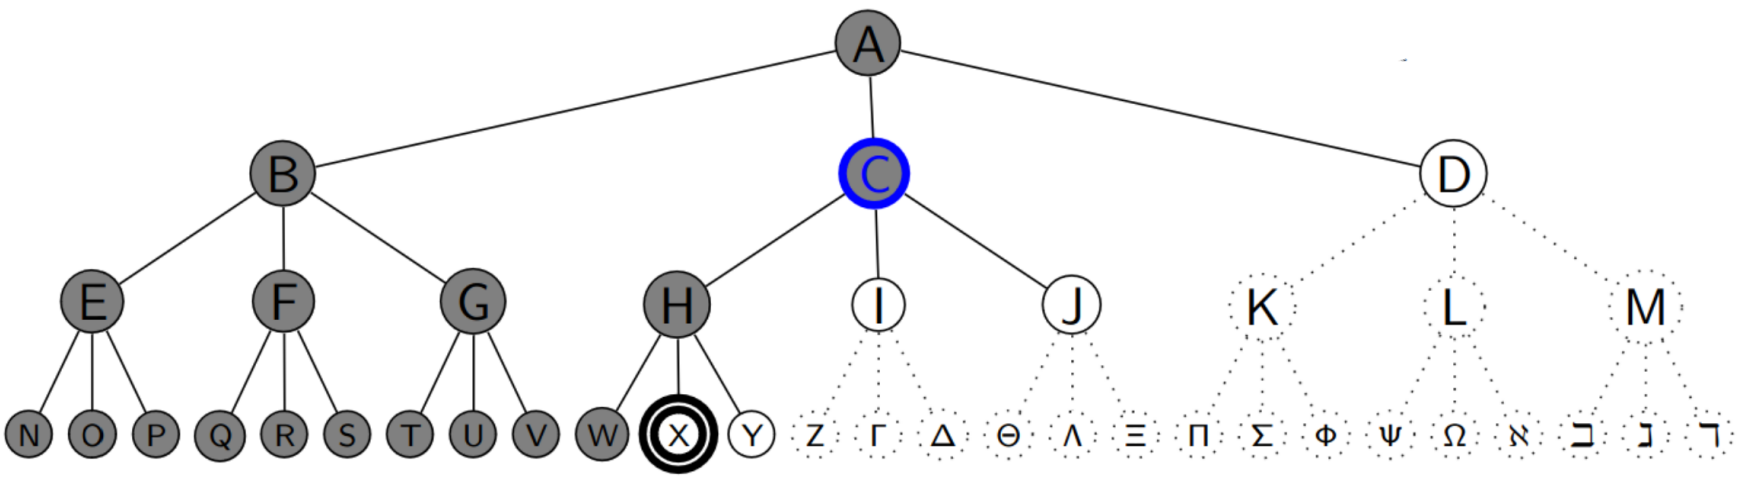

Start of the simulation

- 50s and 60s
- inception of neural networks
- 1956 workshop - describing intelligence precisely
- independent agents in complex environment
- simulating human creativity and language

Problem representation

- how to represent a problem?
- efficiency and computability
- TSP - graph (adjacency matrix)
- how to represent human knowledge, diagnostic data?

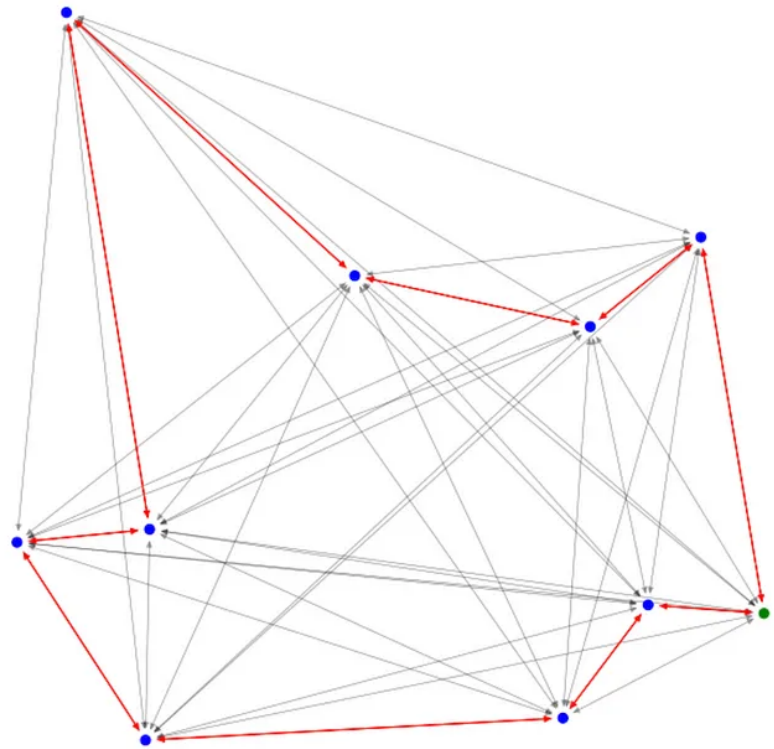

Knowledge bases

Semantic networks

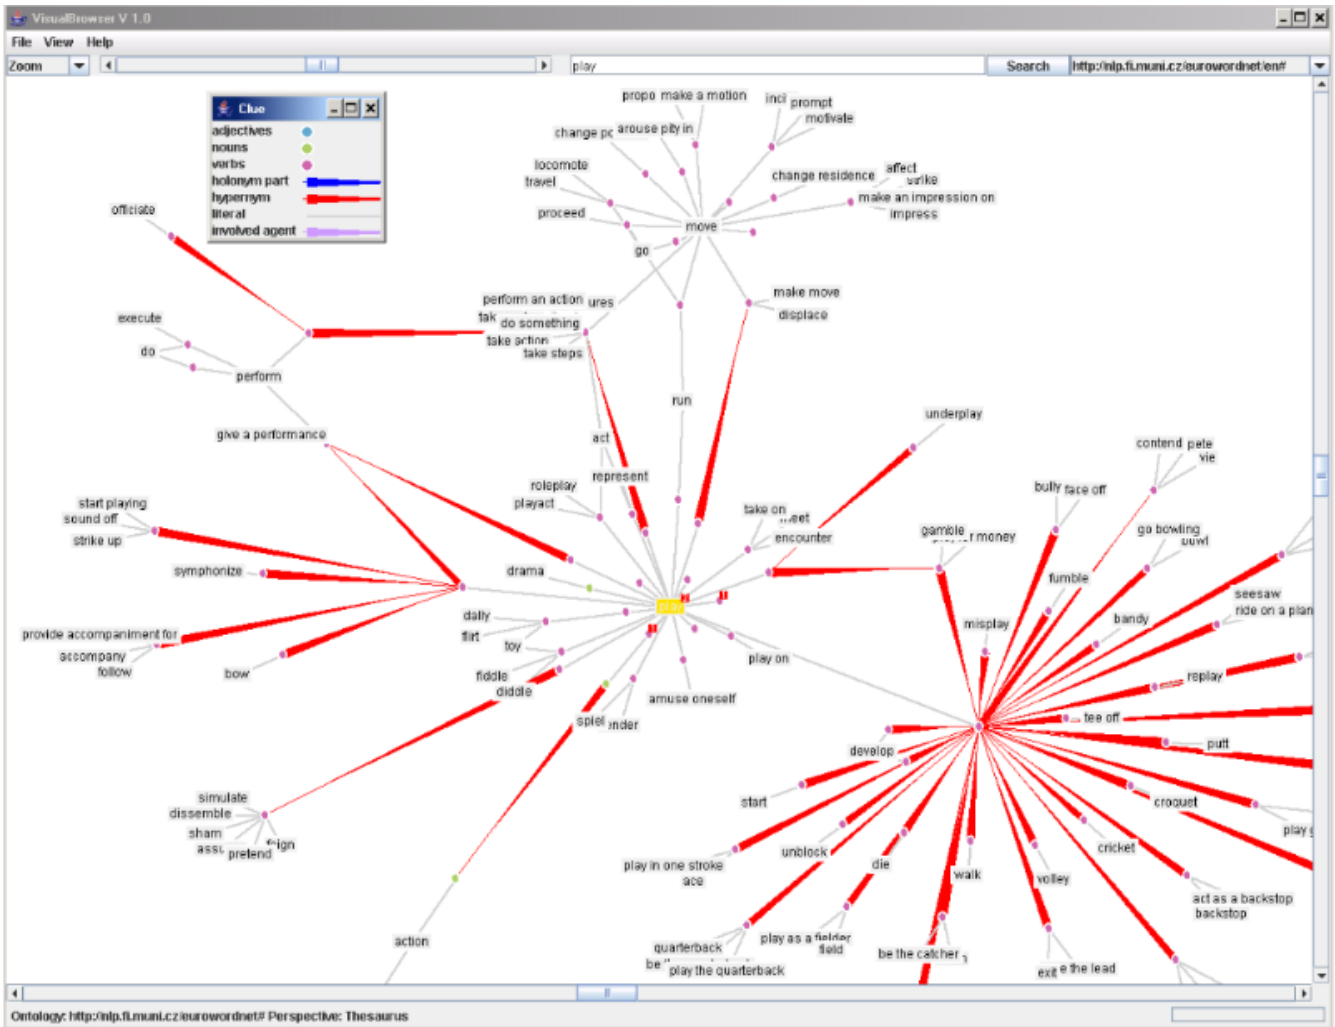

AI Boom$_1$

- Simon, Newell - General Problem Solver
- modeling human reasoning
- 90s, huge investment into AI
- AI in medicine, banking
- companies building 100s of systems

Expert systems

- emulate decision making
- use if-then rules to represent reasoning
- knowledge base to represent knowledge
- inference engines - evaluation functions
- reasoning chains

RULE037

IF: 1) *The identity of the organism is not known with* certainty, and
2) The stain of the organism is gramneg, and
3) The morphology of the organism is rod, and
4) The aerobicity of the organism is aerobic

THEN: There is strongly suggestive evidence (.8) that the class of the organism is enterobacteriaceae

RULE145

IF: 1) The therapy under consideration is one of: cephalothin clindamycin erythromycin lincomycin vancomycin, and
2) Meningitis is an infectious disease diagnosis *for the patient*

THEN: It is definite (1) the therapy under consideration is not a potential therapy for use against the organism

RULE060

IF: *The identity of the organism is bacteroides*

THEN: I recommend therapy chosen from among the following drugs:

1 - clindamycin	(.99)
2 - chloramphenicol	(.99)
3 - erythromycin	(.57)
4 - tetracycline	(.28)
5 - carbenicillin	(.27)

*DeepBlue

- beat Kasparov in a rematch in 1997
- expert system, 200 mil. positions/sec
- evaluation function of 8000 parts
- encoding positions, piece priorities ...

*State space

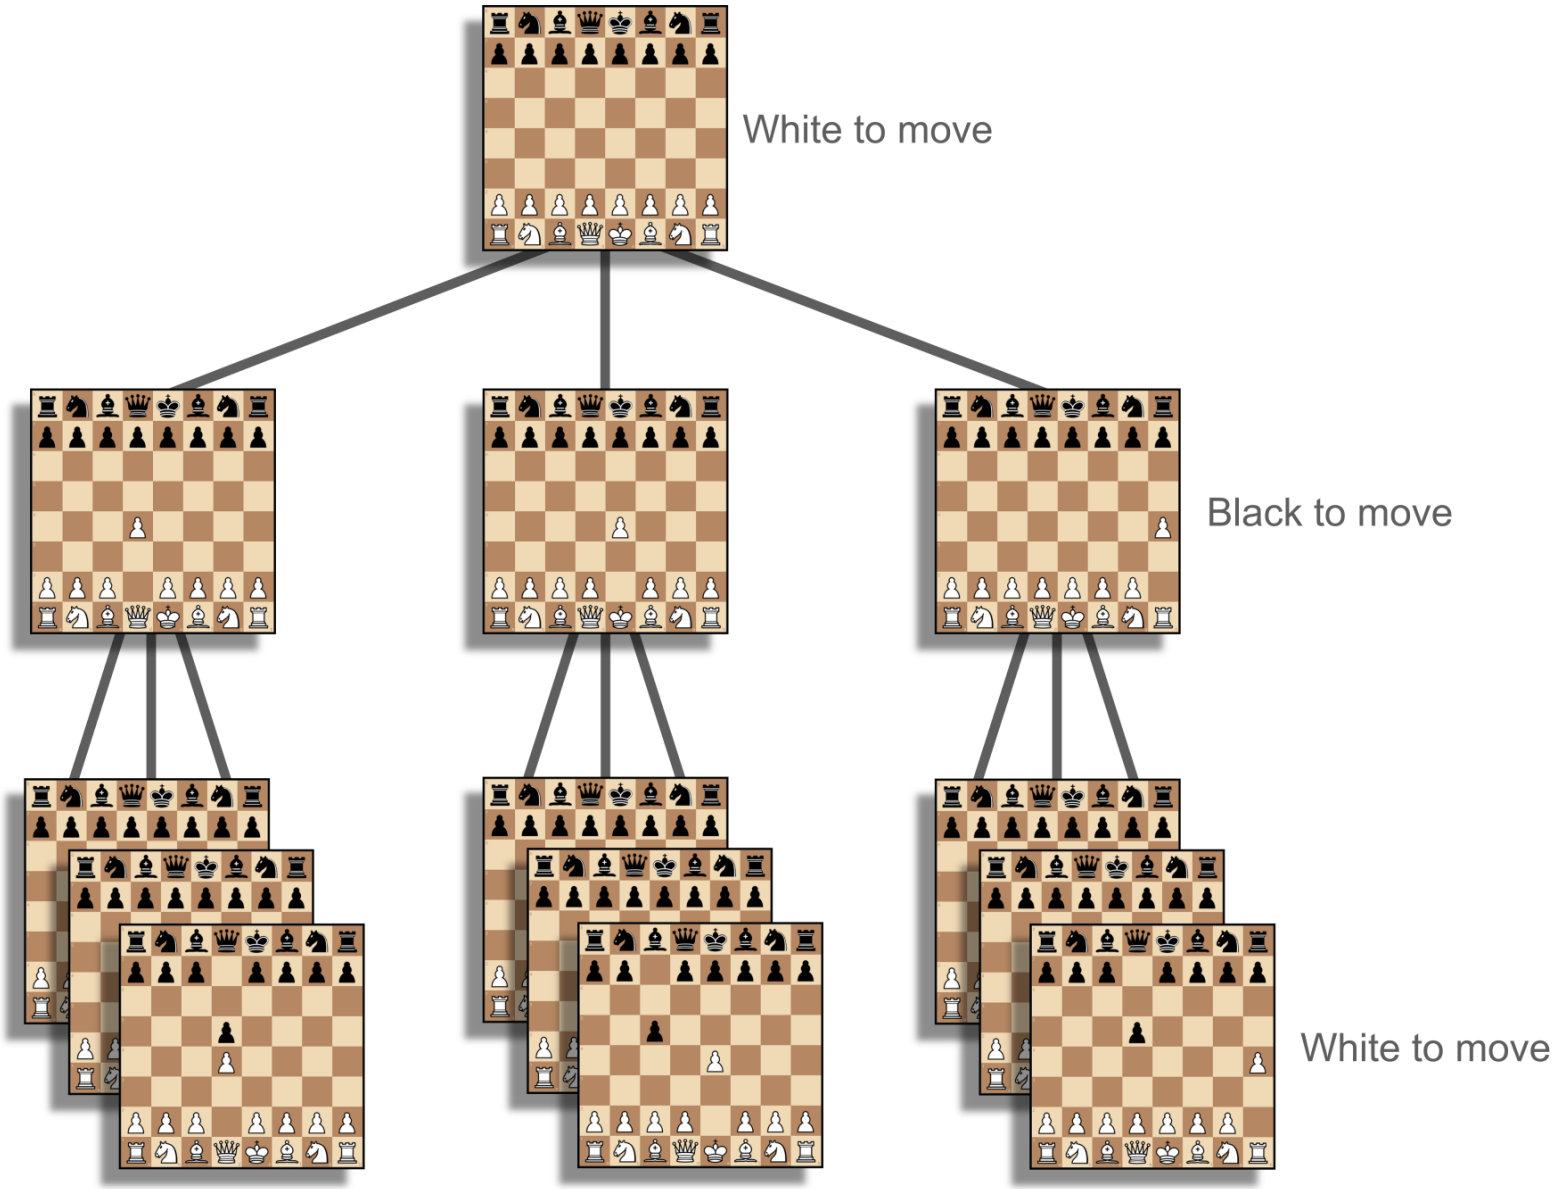

Winter is coming

- expert systems eventually underdelivered
- almost a complete death of the field
- backpropagation for neural networks reinvented
- backprop in cognitive science, cognitive modeling

*Agent modeling

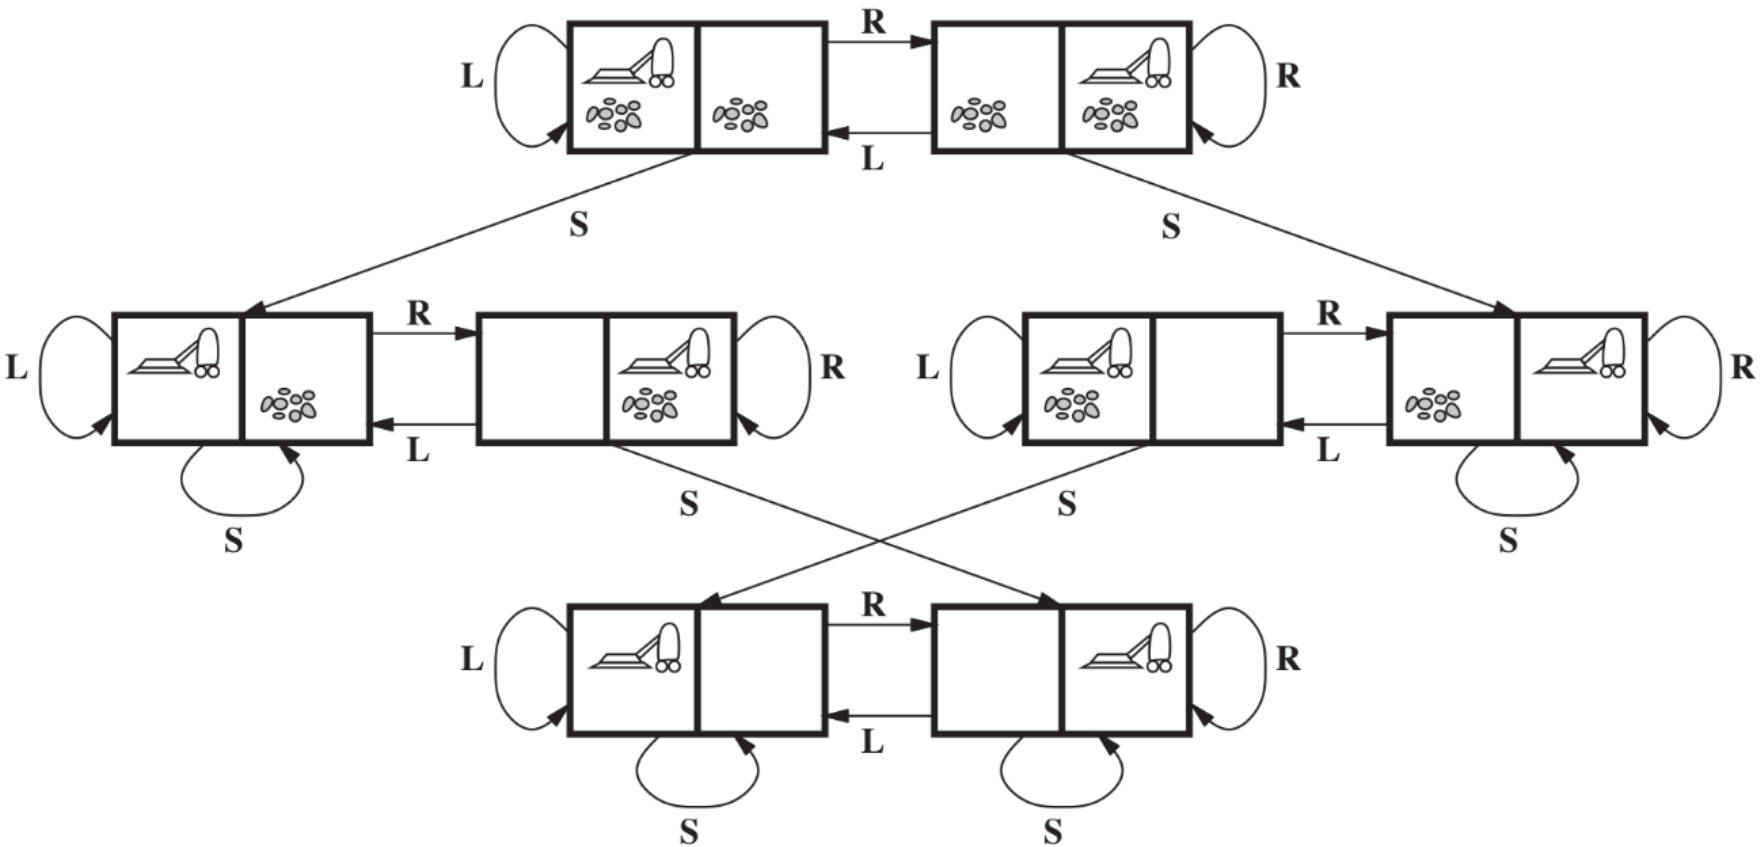

*VLNS and tabu search

- heuristics
- destroy, repair, evaluate
- used in planning
- tabu search keeps recently searched states
- experimentally remembers moves old 5-9 searches

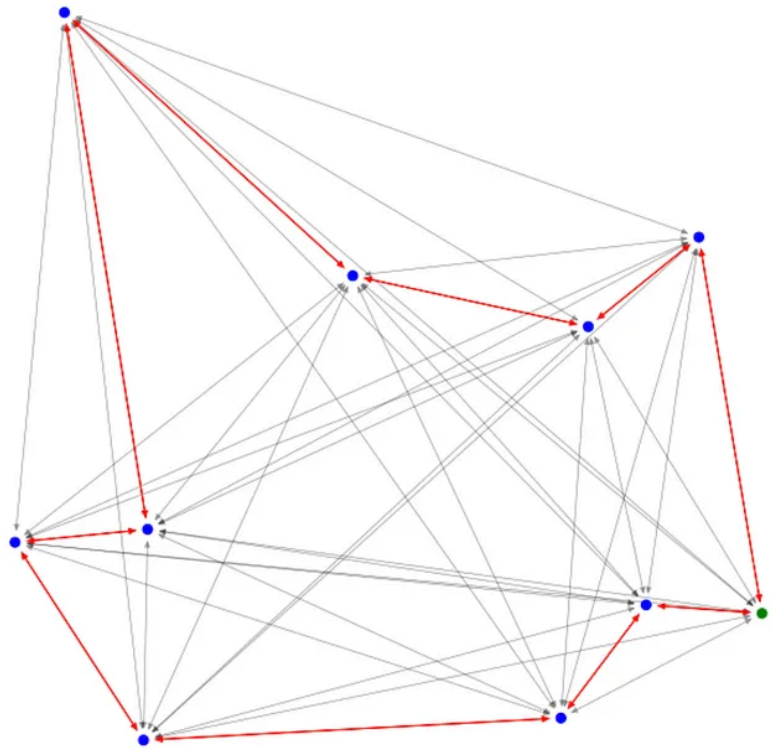

The Spring

- GPUs allowed faster computation
- big data
- machine learning in the forefront
-

Machine learning

- supervised, unsupervised, reinforcement learning
- models "fine-tune" themselves from observed data
- usually to categorize unobserved data
- formally strong models, categorize data well
- evaluation of models necessary

Comparing ML and statistics

- scores over distributions
- testing sets over CI
- prediction over inference
- many features, input variables
- anything (data-wise) goes

Linear regression, an ML approach

- $\hat{y}_i = a + bx = w_1 + w_2 x$
- not how it's done in statistics
- a dataset is a set of tuples
- $\{(x_1, y_1), (x_2, y_2), (x_3, y_3), \ldots\}$
- square error - $\sum_{i=1}^{n} (y_i - \hat{y}_i)^2$
- going after minimal error

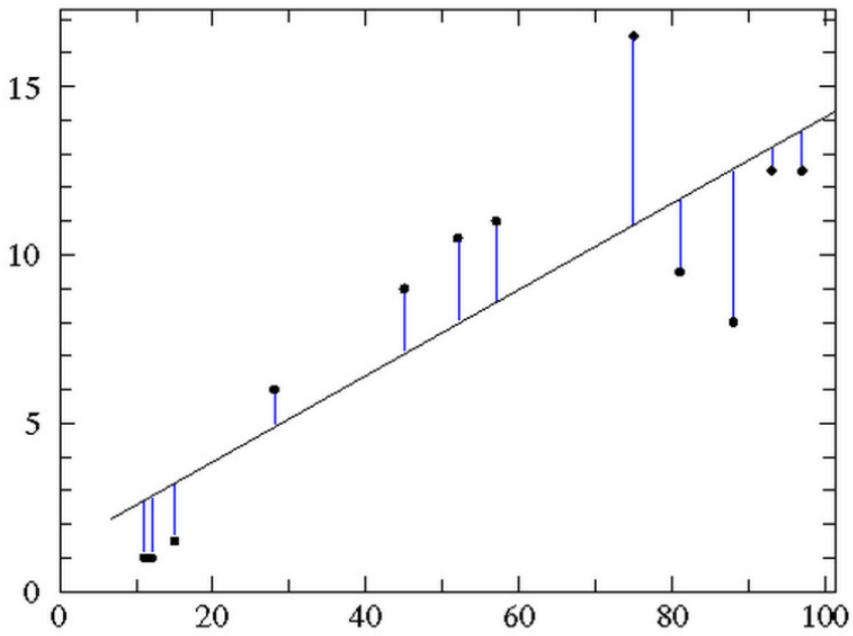

*Gradient descent

- square error for a single point
- $(y_i - \hat{y}_i)^2$
- gradient of square error for a single point
- state space

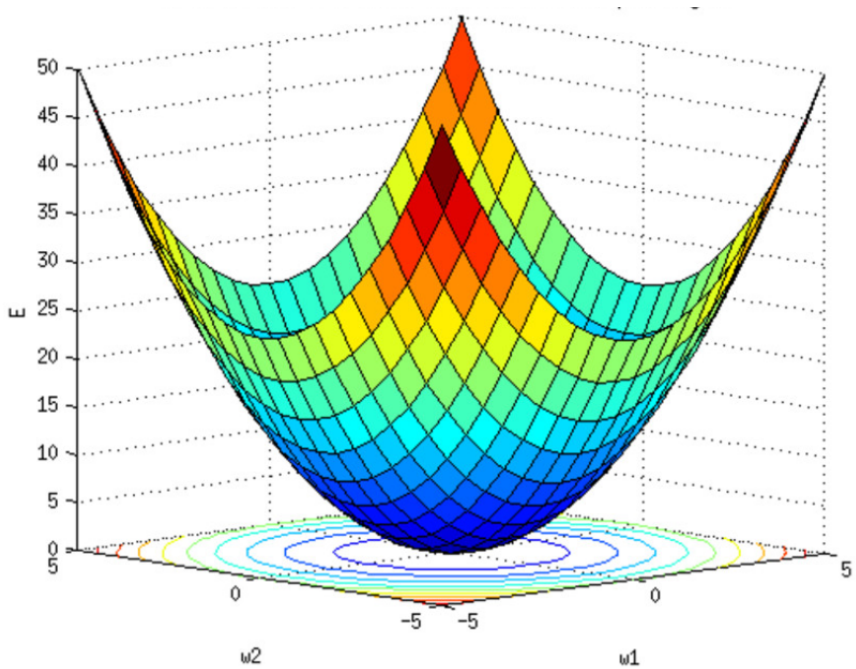

Neural networks

Biological to artificial neuron

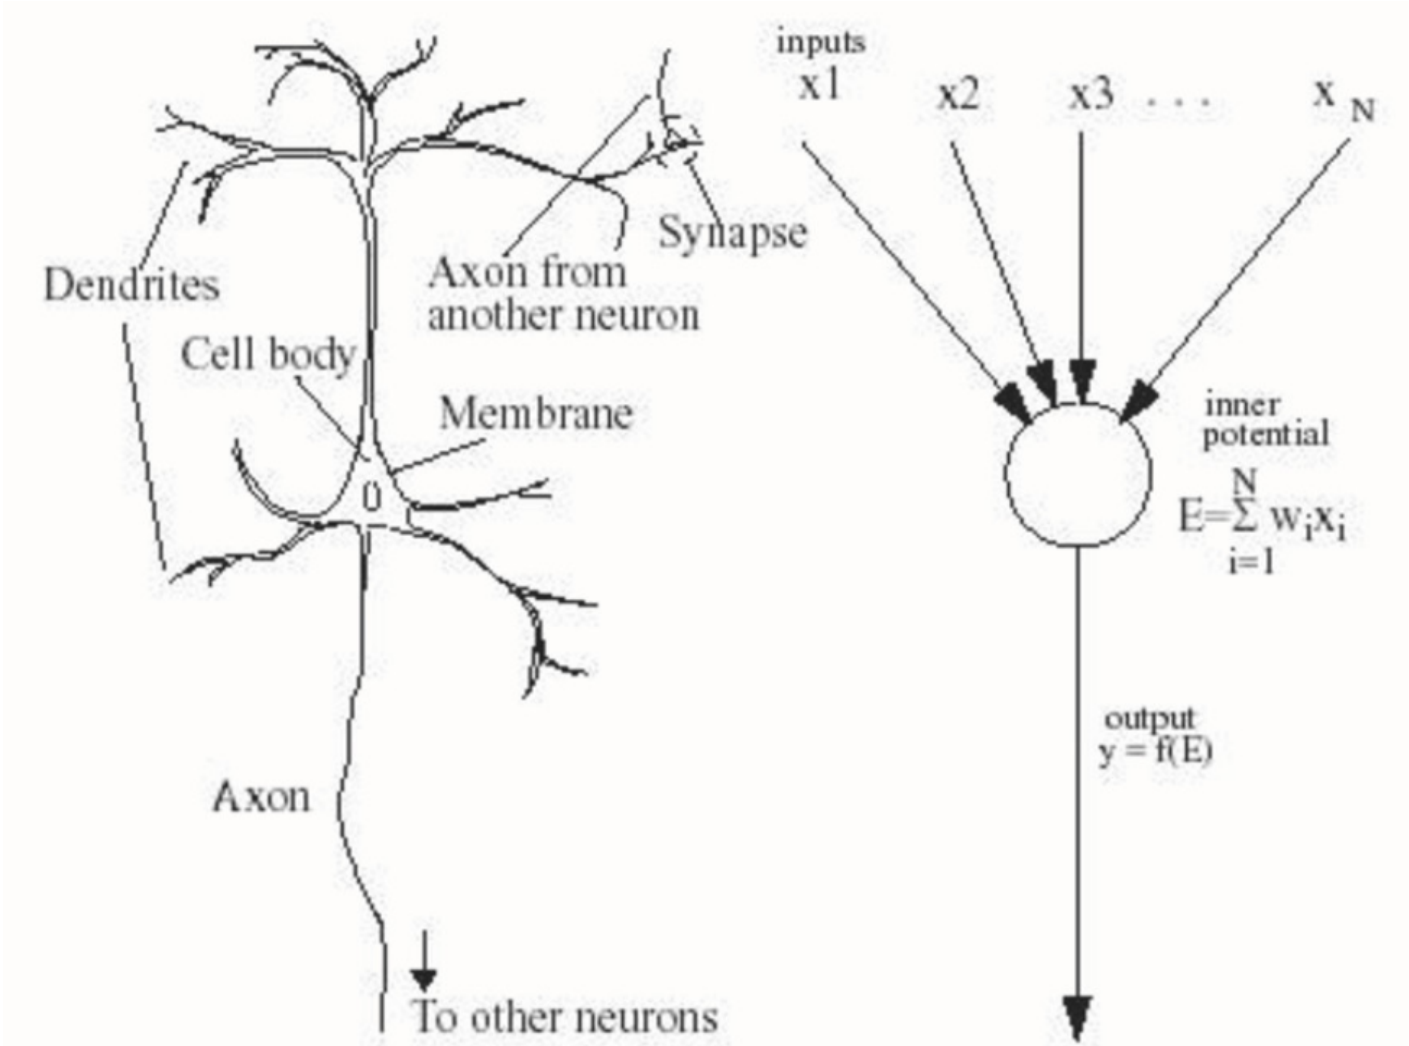

Learning and gradient descent

- square error
- going after minimal error
- compute the direction - gradient
- go opposite to it

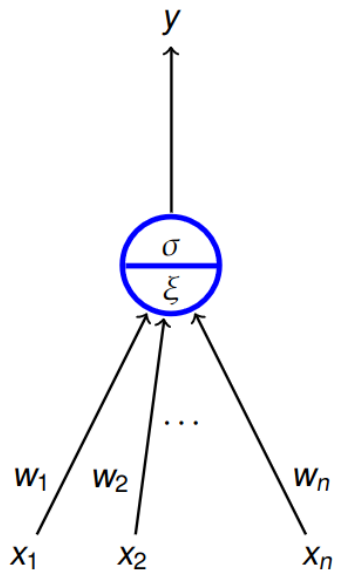

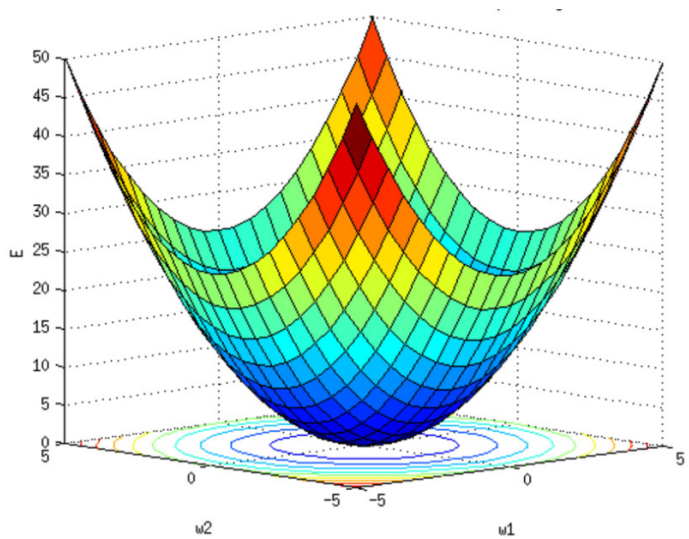

Multi-layer NN

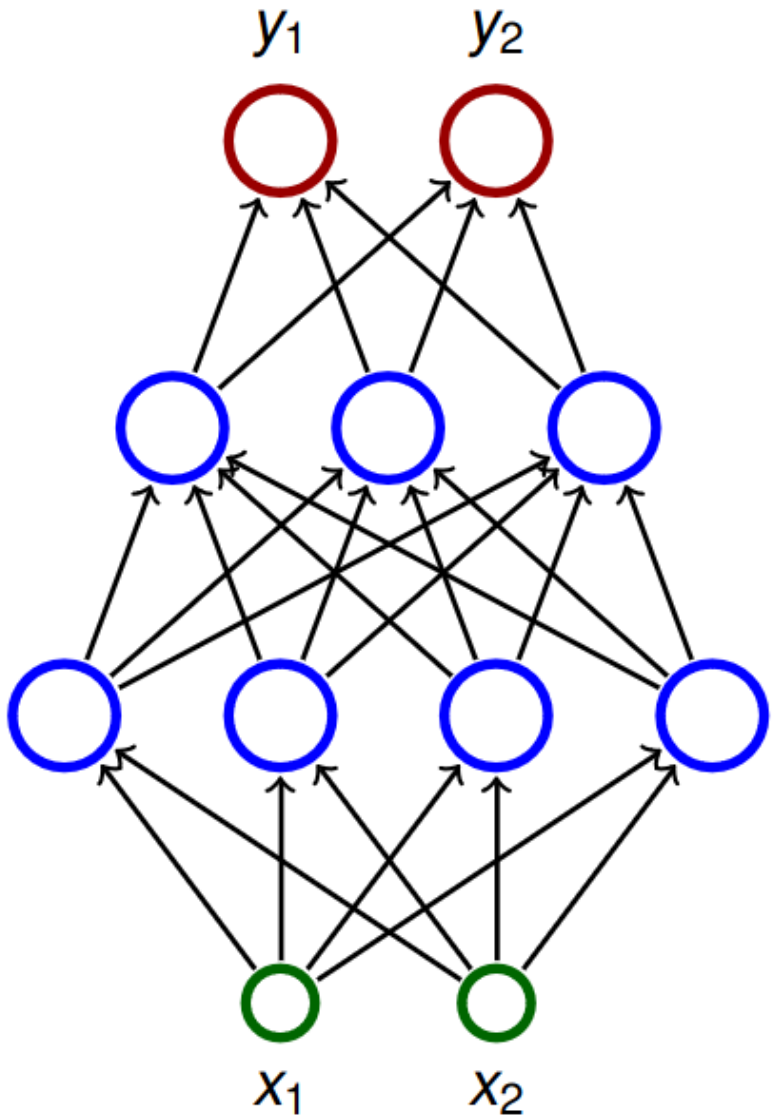

Learning in deeper NNs

- square error for output layer
- backprop for the rest

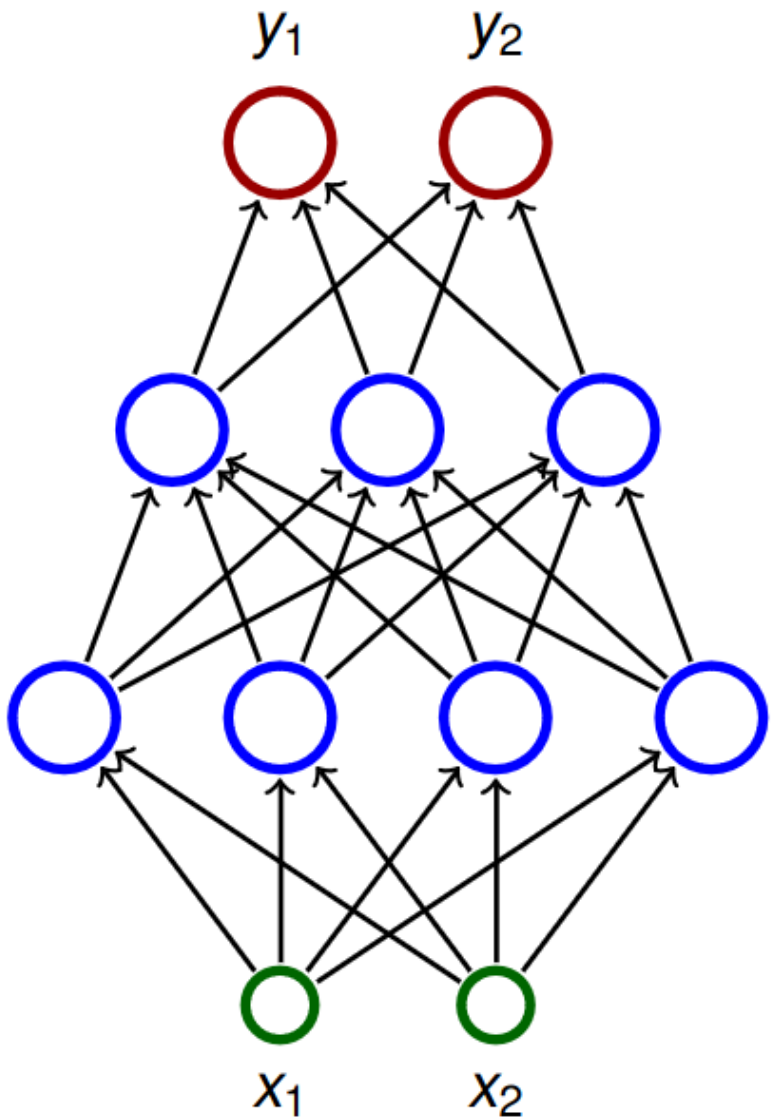

*Activation functions

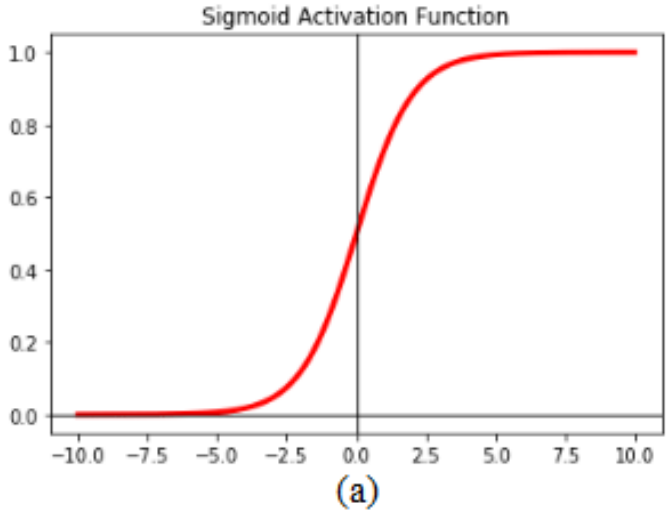

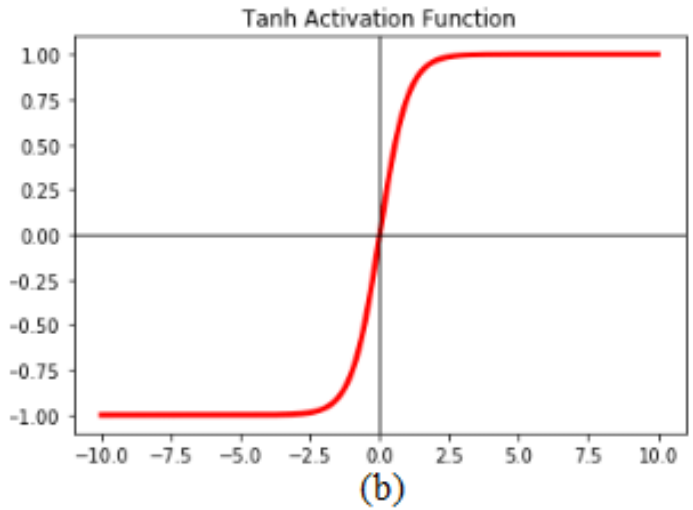

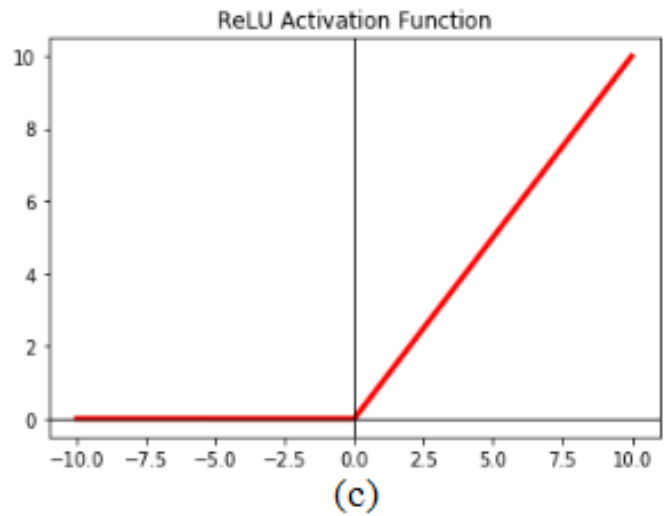

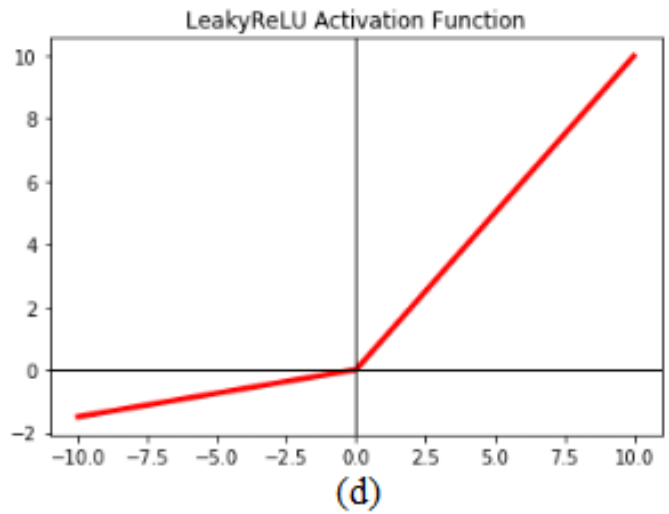

AI boom$_2$

- Convolutional NNs, Reccurent NNs, Self-attention
- generalizability
- interpretability and explainability
- Explainable AI (XAI)
- Data-Centric AI (DCAI)

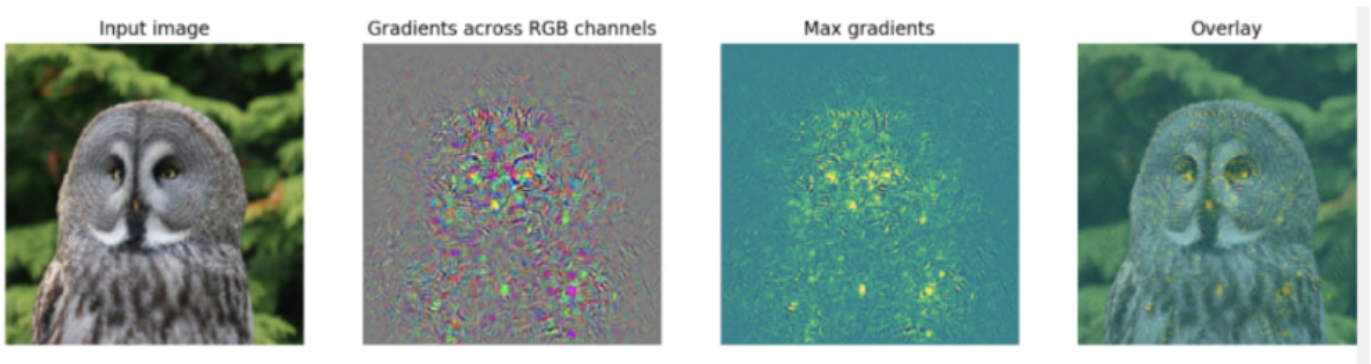

MASARYK
UNIVERSITY

Sources

- Abraham, A. (2005). Rule-Based expert systems. Handbook of measuring system design.
- Campbell, M., Hoane Jr, A. J., Hsu, F. H. (2002). Deep blue. Artificial
- Russell, S. J., & Norvig, P. (2010). Artificial intelligence a modern approach. London. intelligence, 134(1-2), 57-83.
- Kandel, I., Castelli, M. (2020). Transfer learning with convolutional neural networks for diabetic retinopathy image classification. A review. Applied Sciences, 10(6), 2021.
- Slides of PB016, Aleš Horák
- Slides of IB031, PV021, Tomáš Brázdil
- https://chrisbutner.github.io/ChessCoach/high-level-explanation.html
- https://www.flickr.com/photos/jamesthephotographer/