AlphaGo

Как думает машина

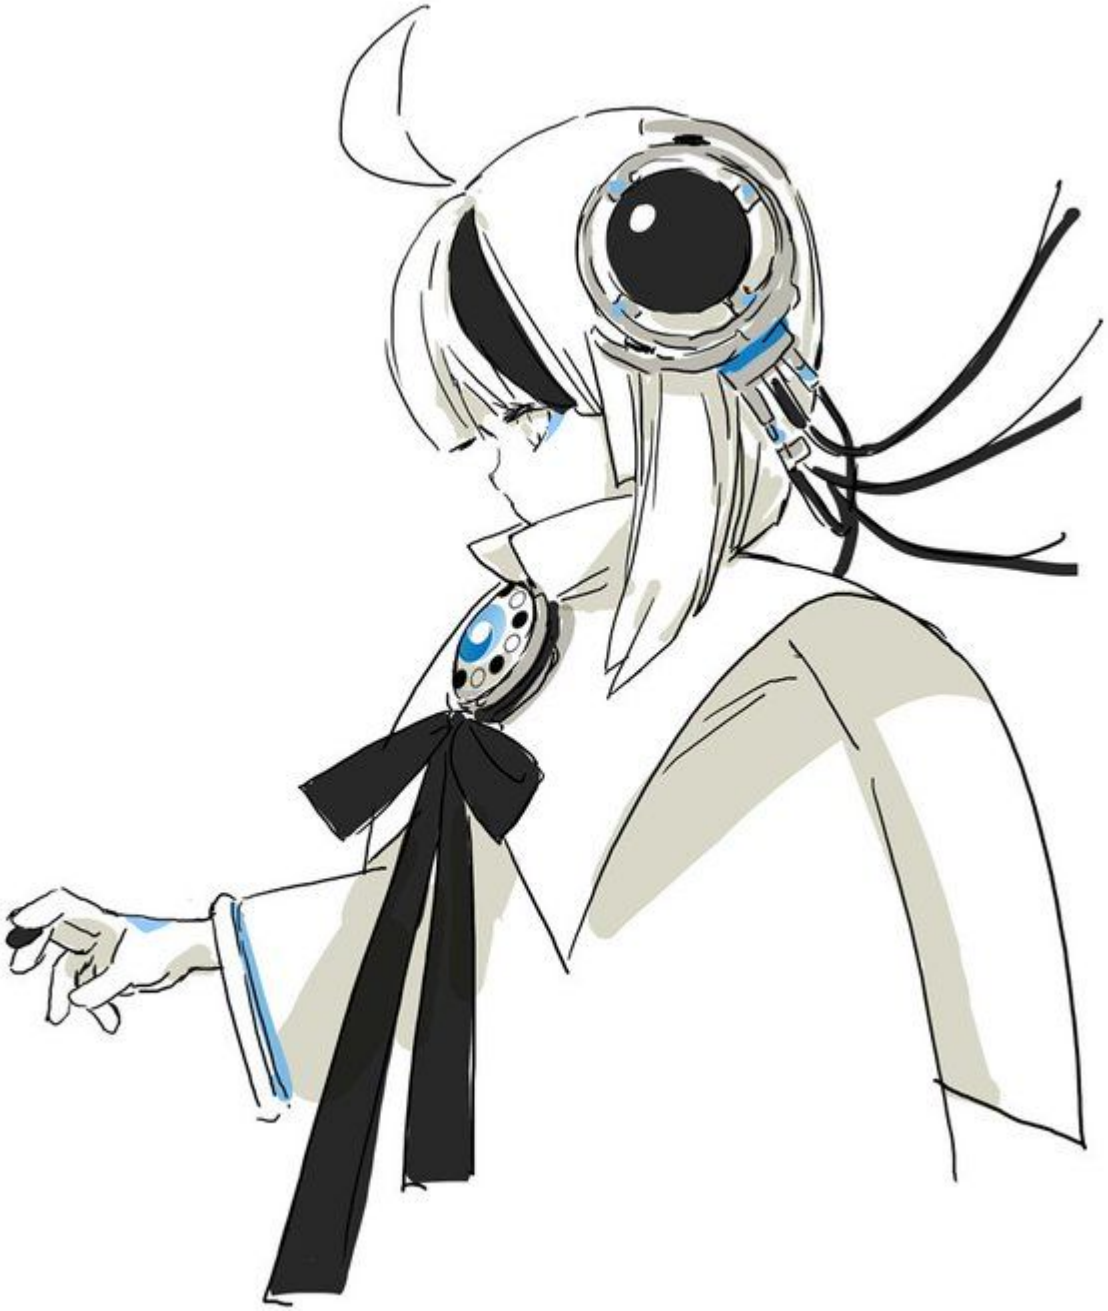

Google DeepMind
Challenge Match
8 - 15 March 2016
AlphaGo
ALPHAGO
02:00:00
LEE SEDOL
02:00:00
AlphaGo
Lee Sedol

В чем интерес

- Исторический контекст
- Подходы к решению задач
- Внутри AlphaGo

Компьютеры vs Человеки

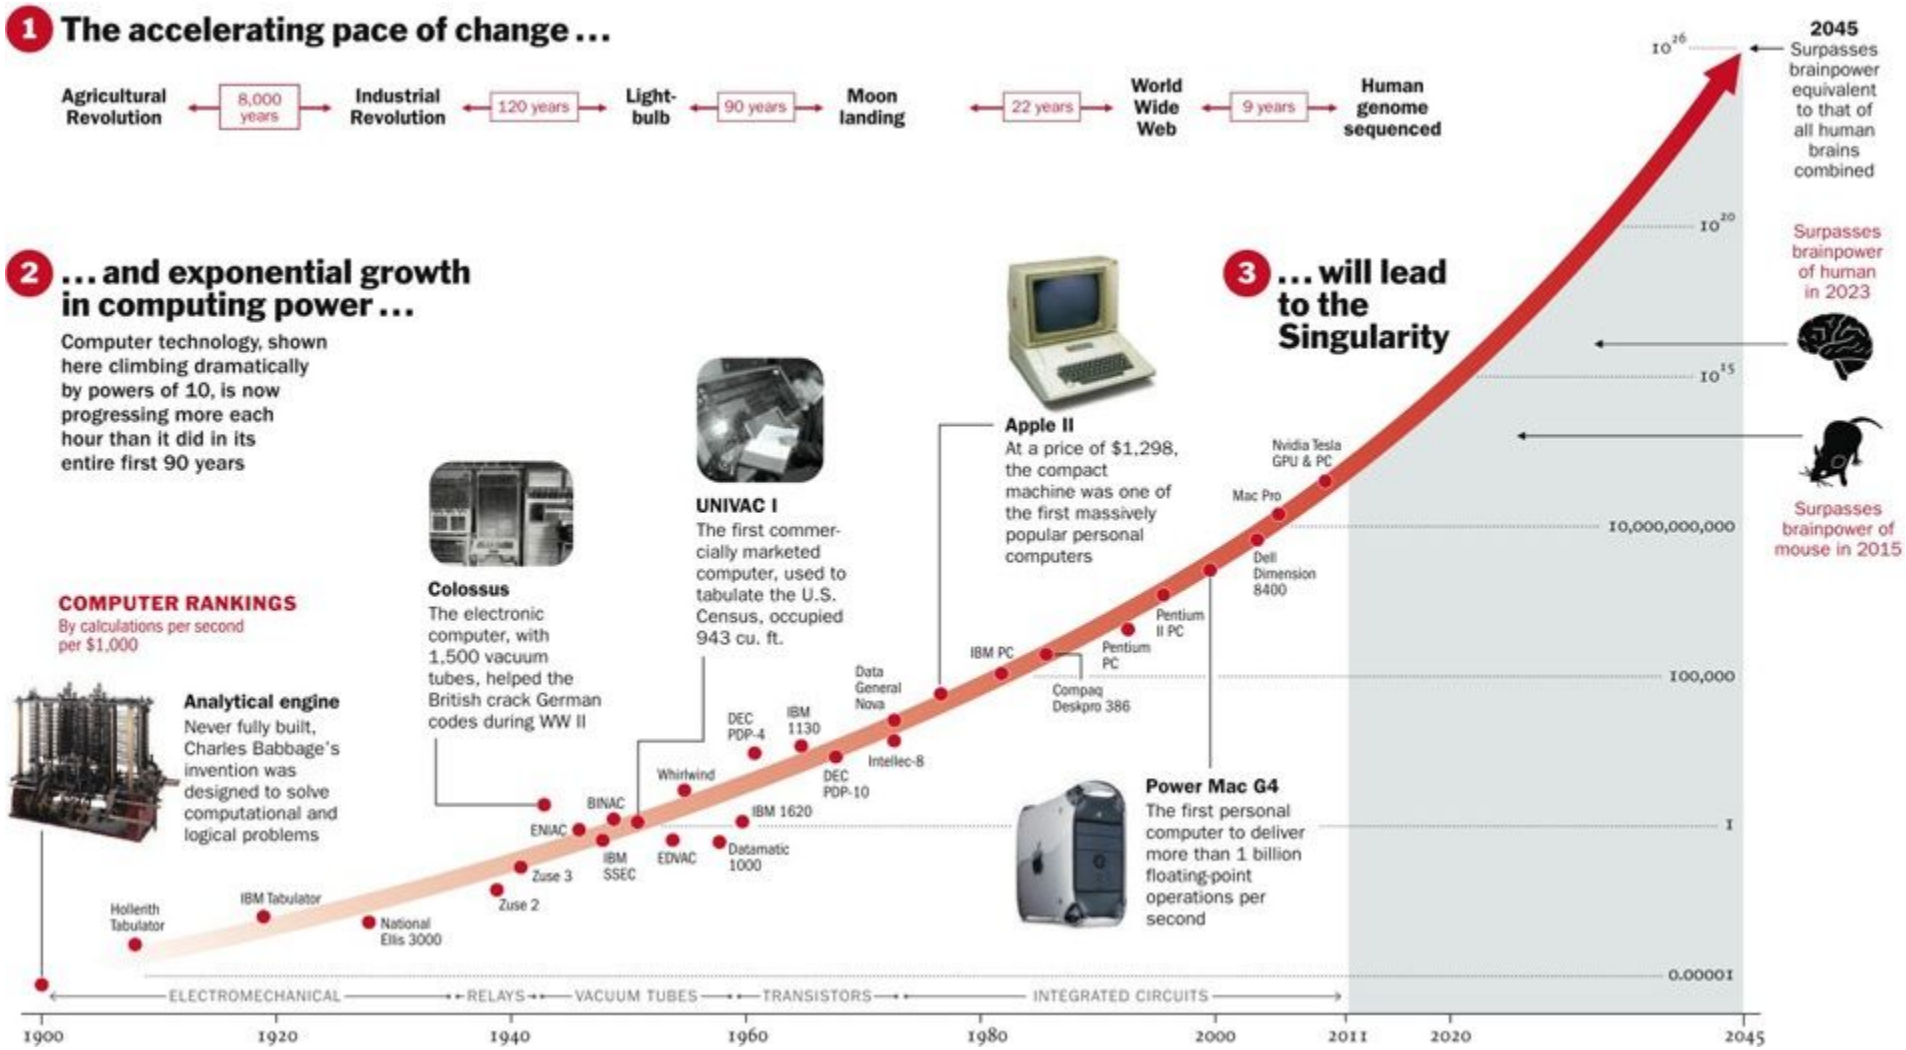

В ИНФОРМАТИКЕ БЫВАЕТ ТРУДНО ОБЪЯСНИТЬ РАЗНИЦУ МЕЖДУ ПРОСТЫМ И ПРАКТИЧЕСКИ НЕВОЗМОЖНЫМ.

Способ решения задач

- Грубая сила — перебор
- Умная сила — тайные знания
  - вручную (алгоритмы)
  - автомагически
    - машинное обучение
    - нейронные сети

Нейронная сеть — Человек

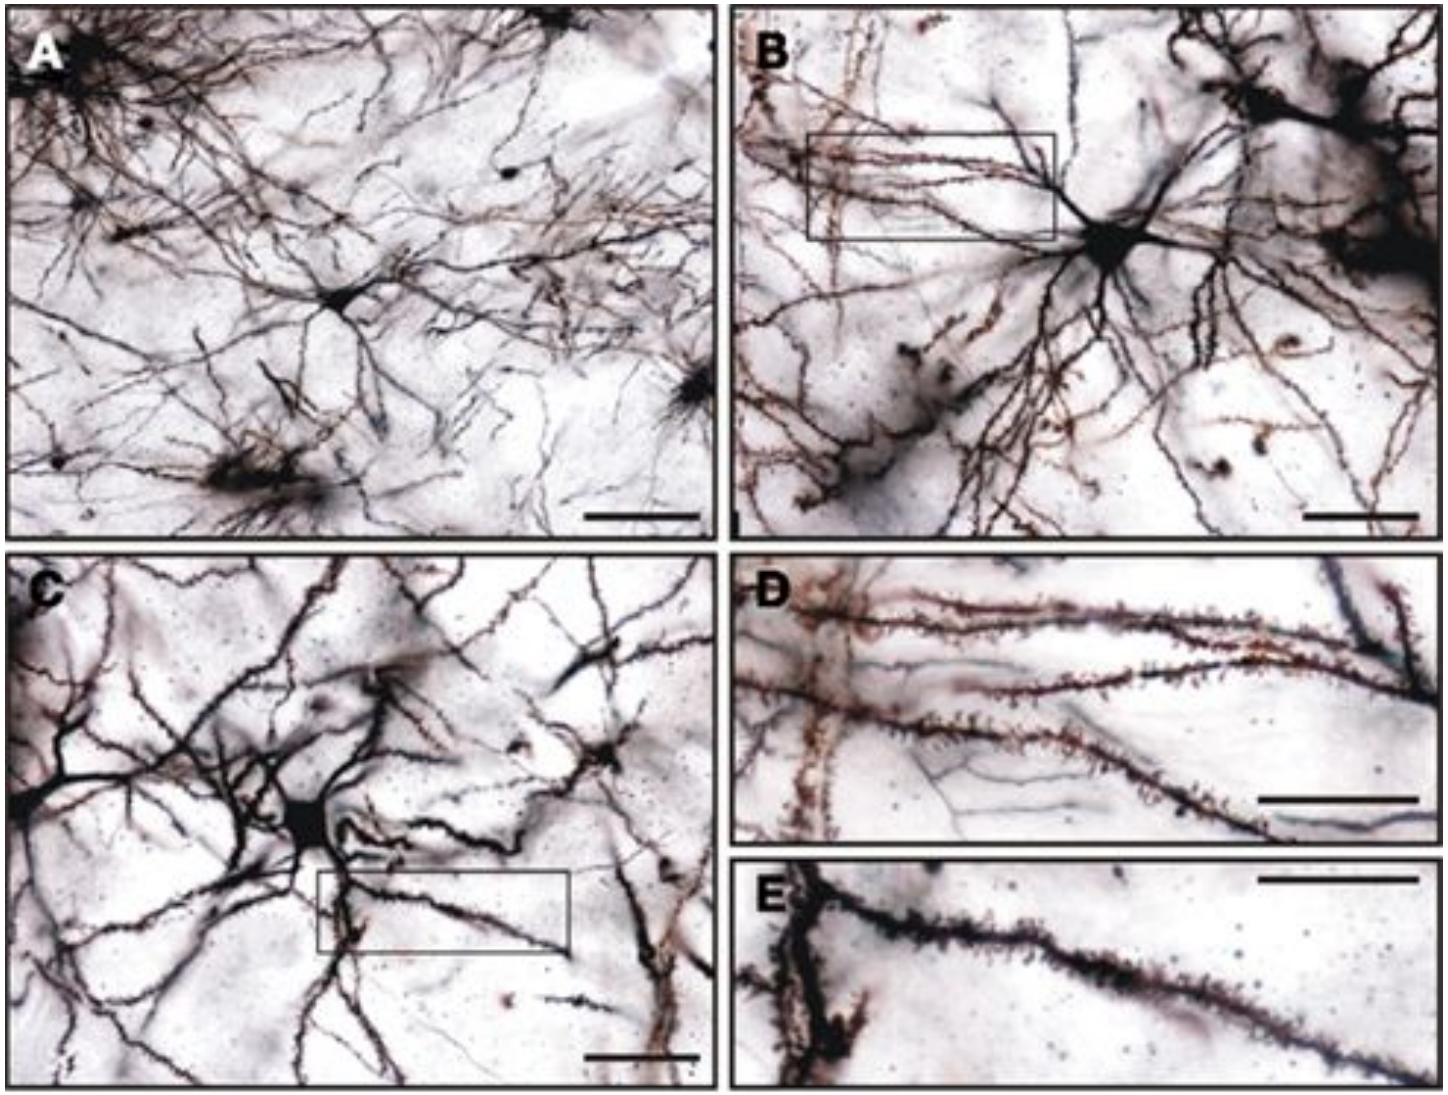

Нейронная сеть — Программа

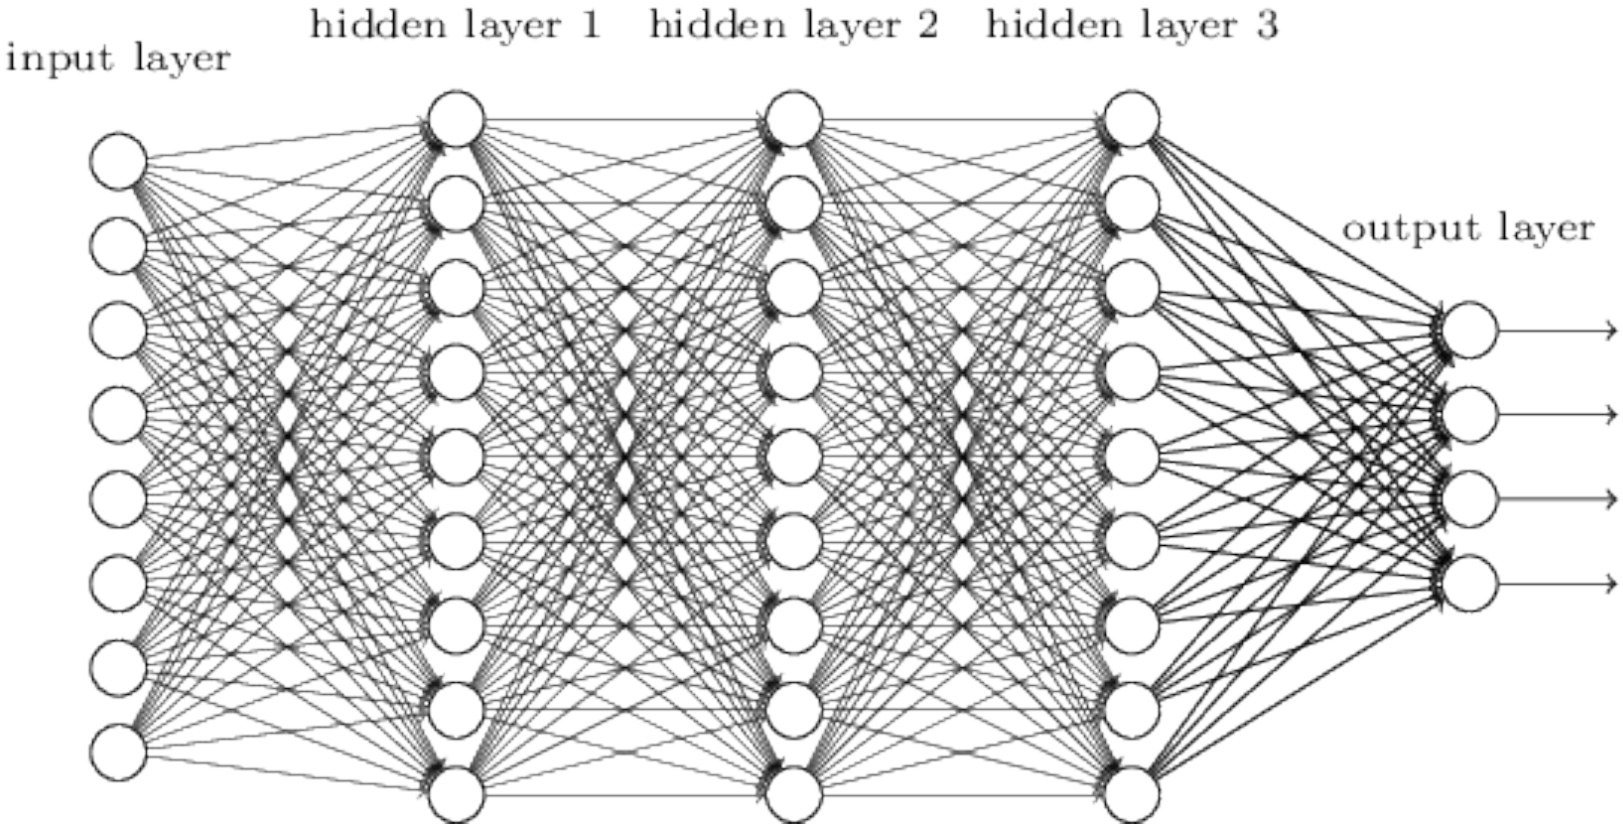

Свёрточная глубокая нейронная сеть

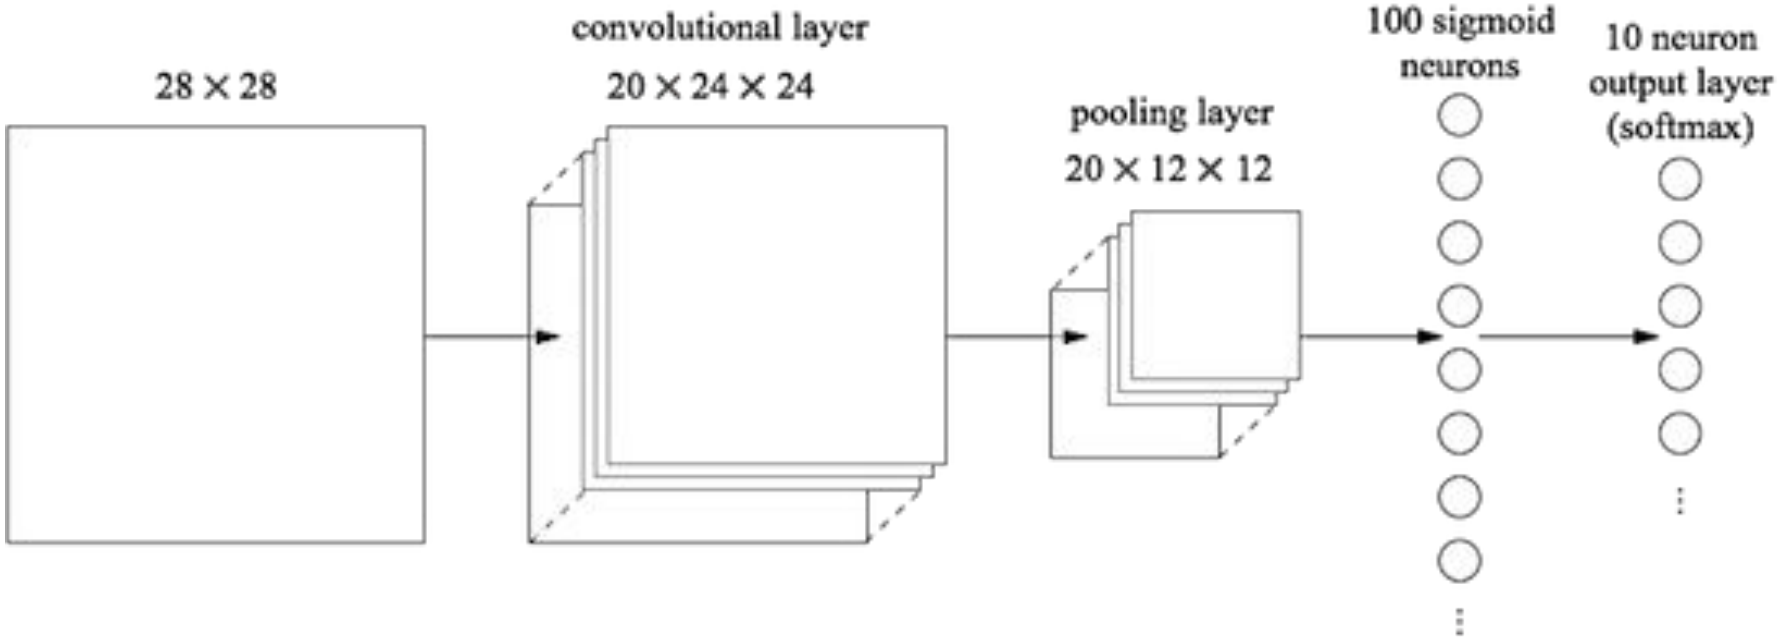

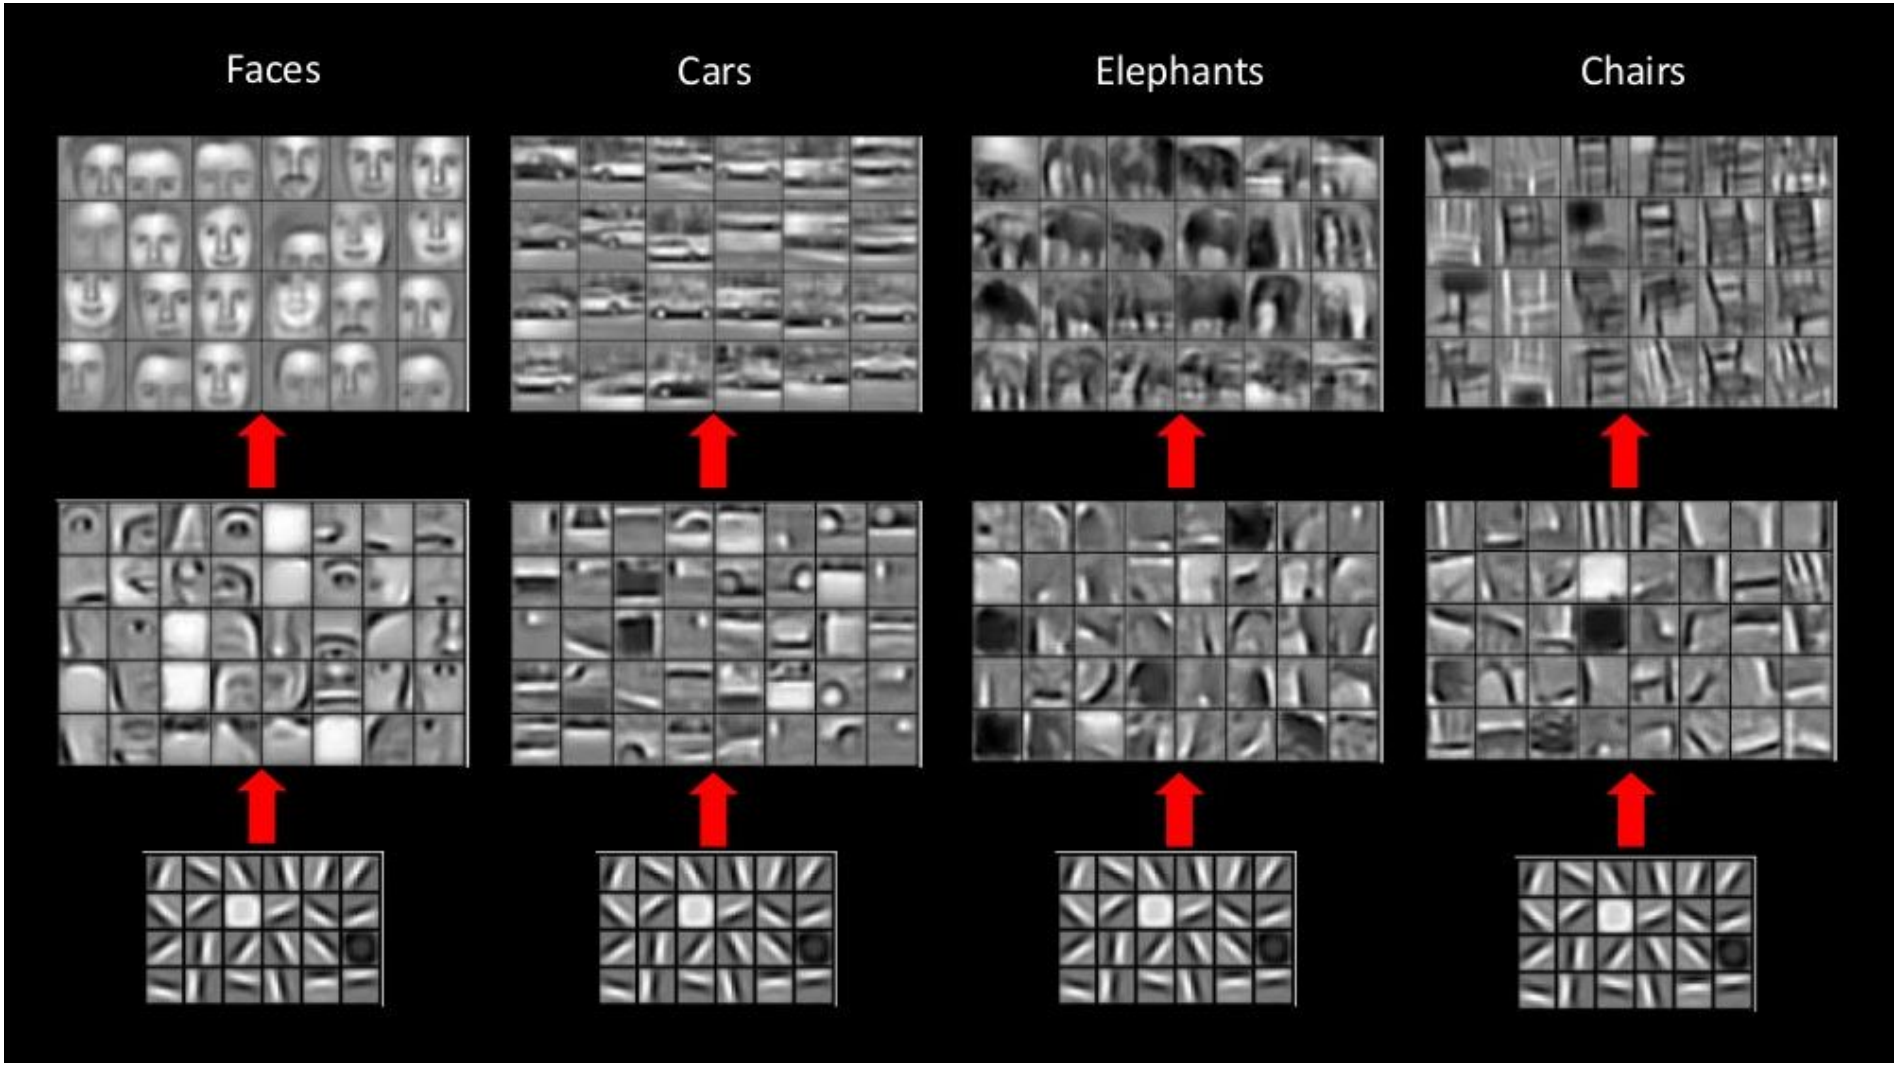
Faces
Cars
Elephants
Chairs

WHAT MAKES DEEP LEARNING DEEP?

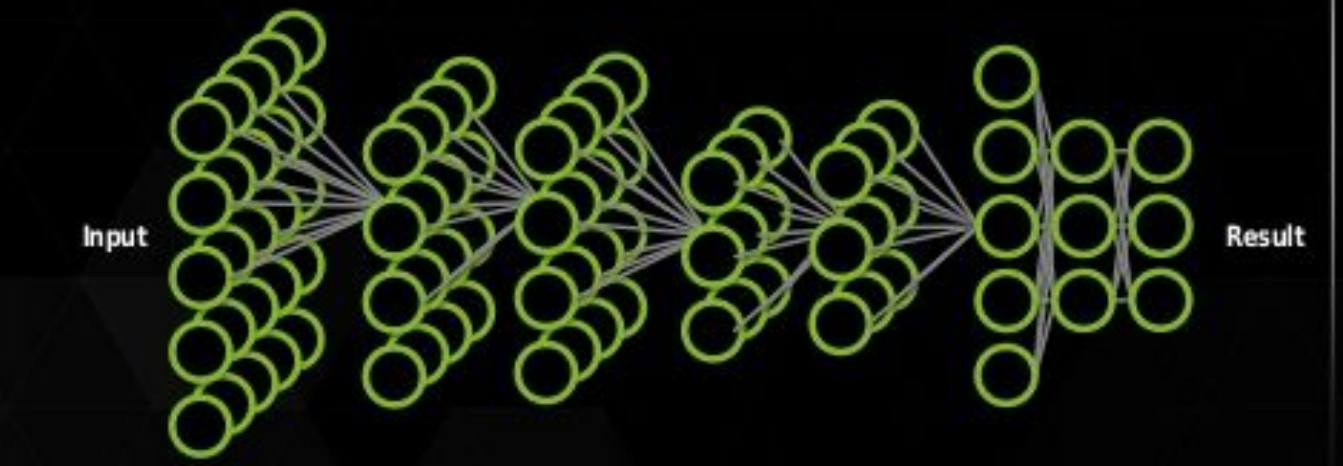

Today's Largest Networks

~10 layers
1B parameters
10M images
~30 Exaflops
~30 GPU days

Human brain has trillions of parameters - only 1,000 more.

MACHINE LEARNING USE CASES

...machine learning is pervasive

Image Classification, Object Detection, Localization

Face Recognition

Speech & Natural Language Processing

Medical Imaging & Interpretation

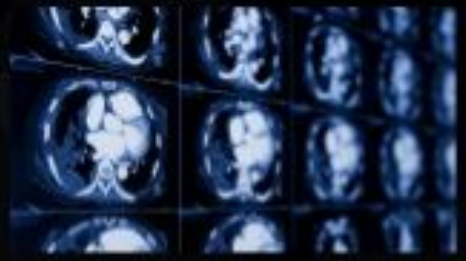

Seismic Imaging & Interpretation

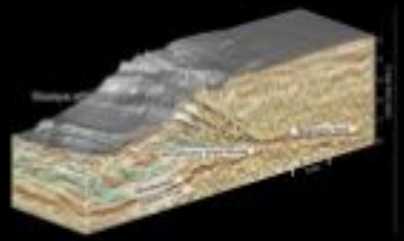

Recommendation

Внутри DeepBlue (1997)

- 200,000,000 позиций в секунду
- Материал, позиция, сохранность Короля, темп
- 11.38 GFLOPS
- Специализированное железо

Go

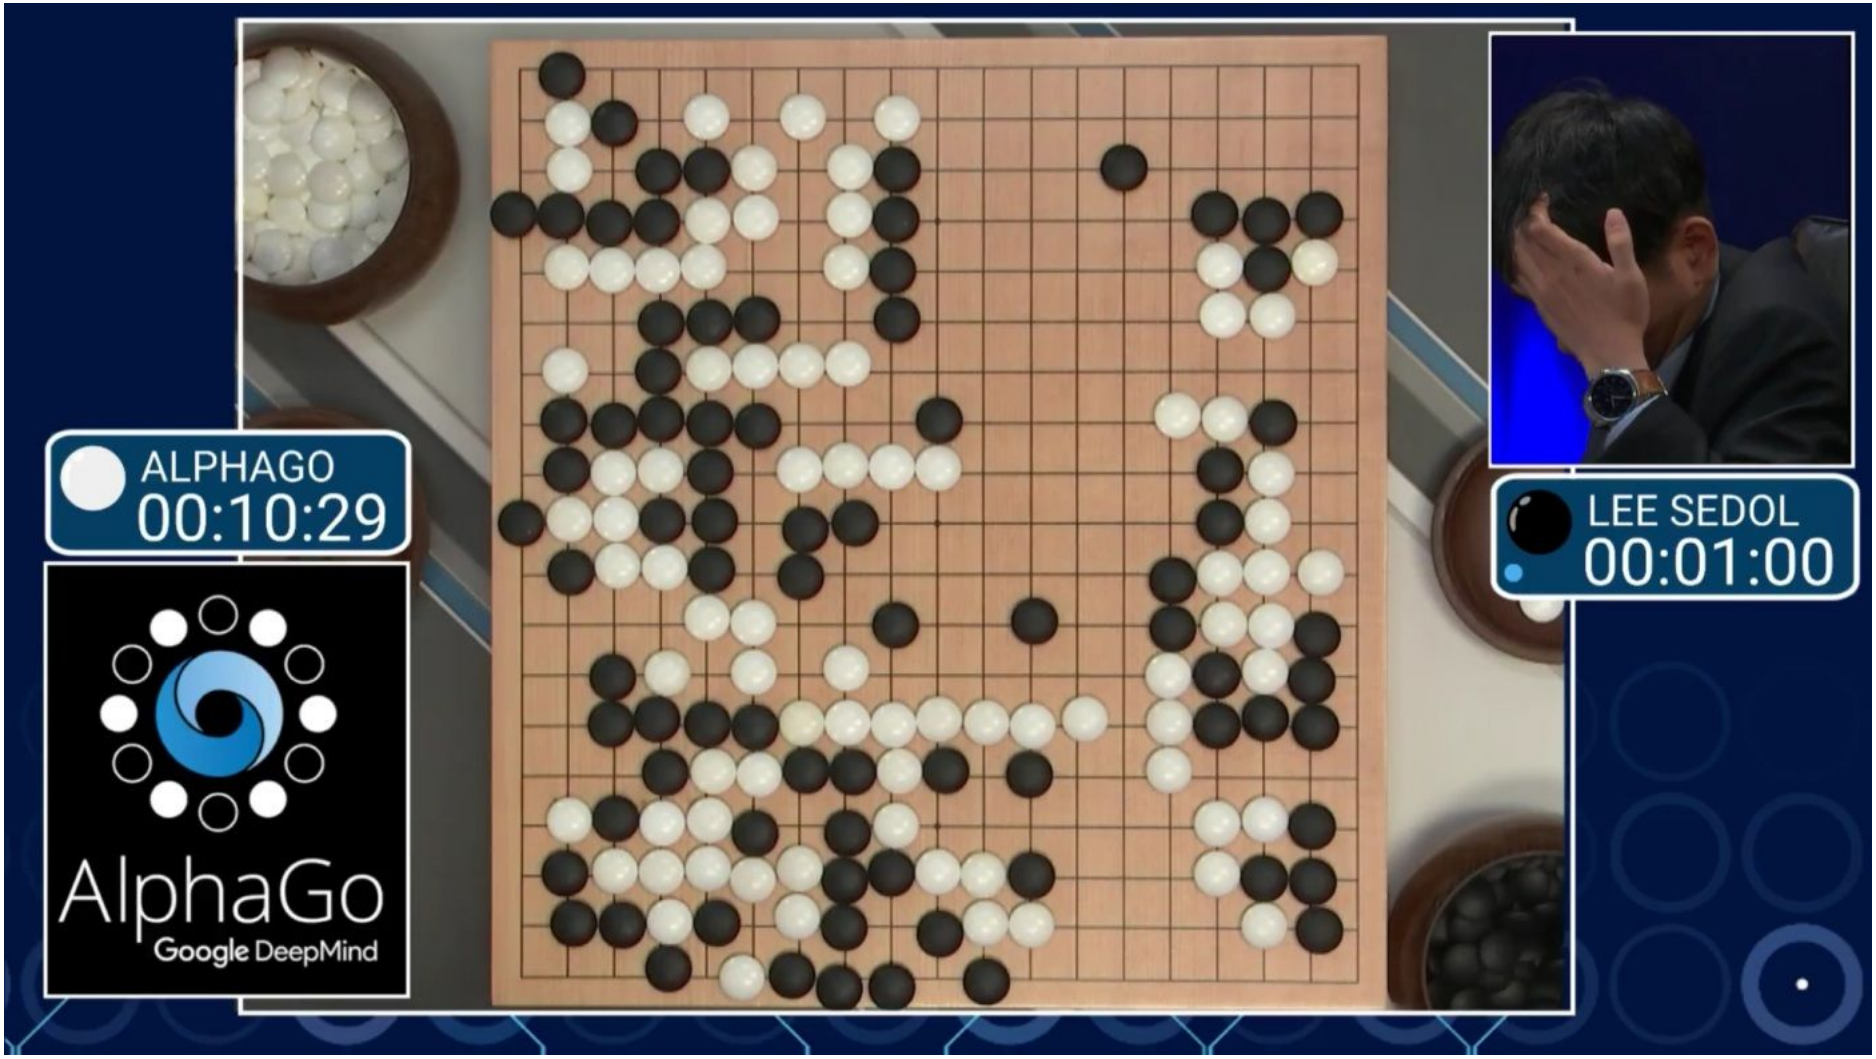

Внутри Alpha Go

- Features
  - 46 bits for slow
  - ≈140,000 patterns for fast
- Value networks to evaluate board positions
- Policy networks to select moves
- Fast value network for rollout
- Monte-Carlo tree search
- “machine learning techniques” for time management

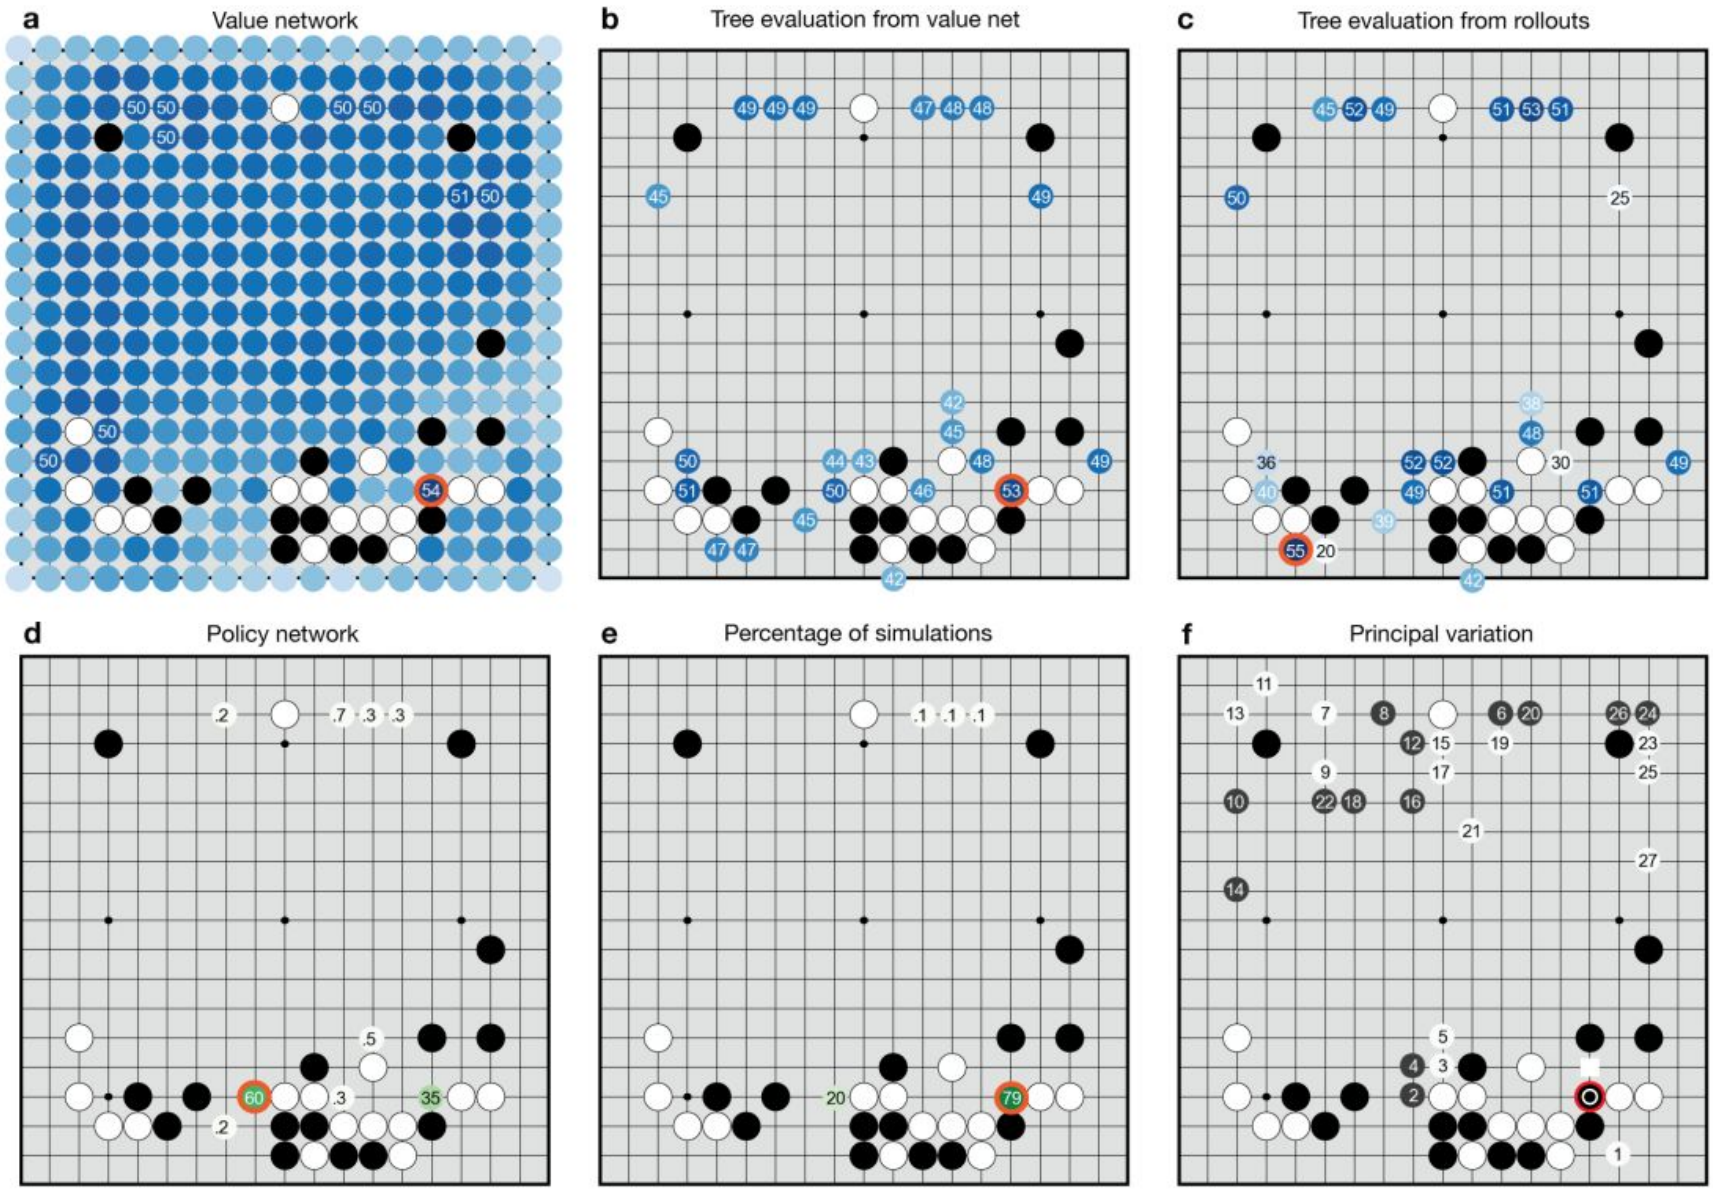
a
Value network
b
Tree evaluation from value net
c
Tree evaluation from rollouts
d
Policy network
e
Percentage of simulations
f
Principal variation

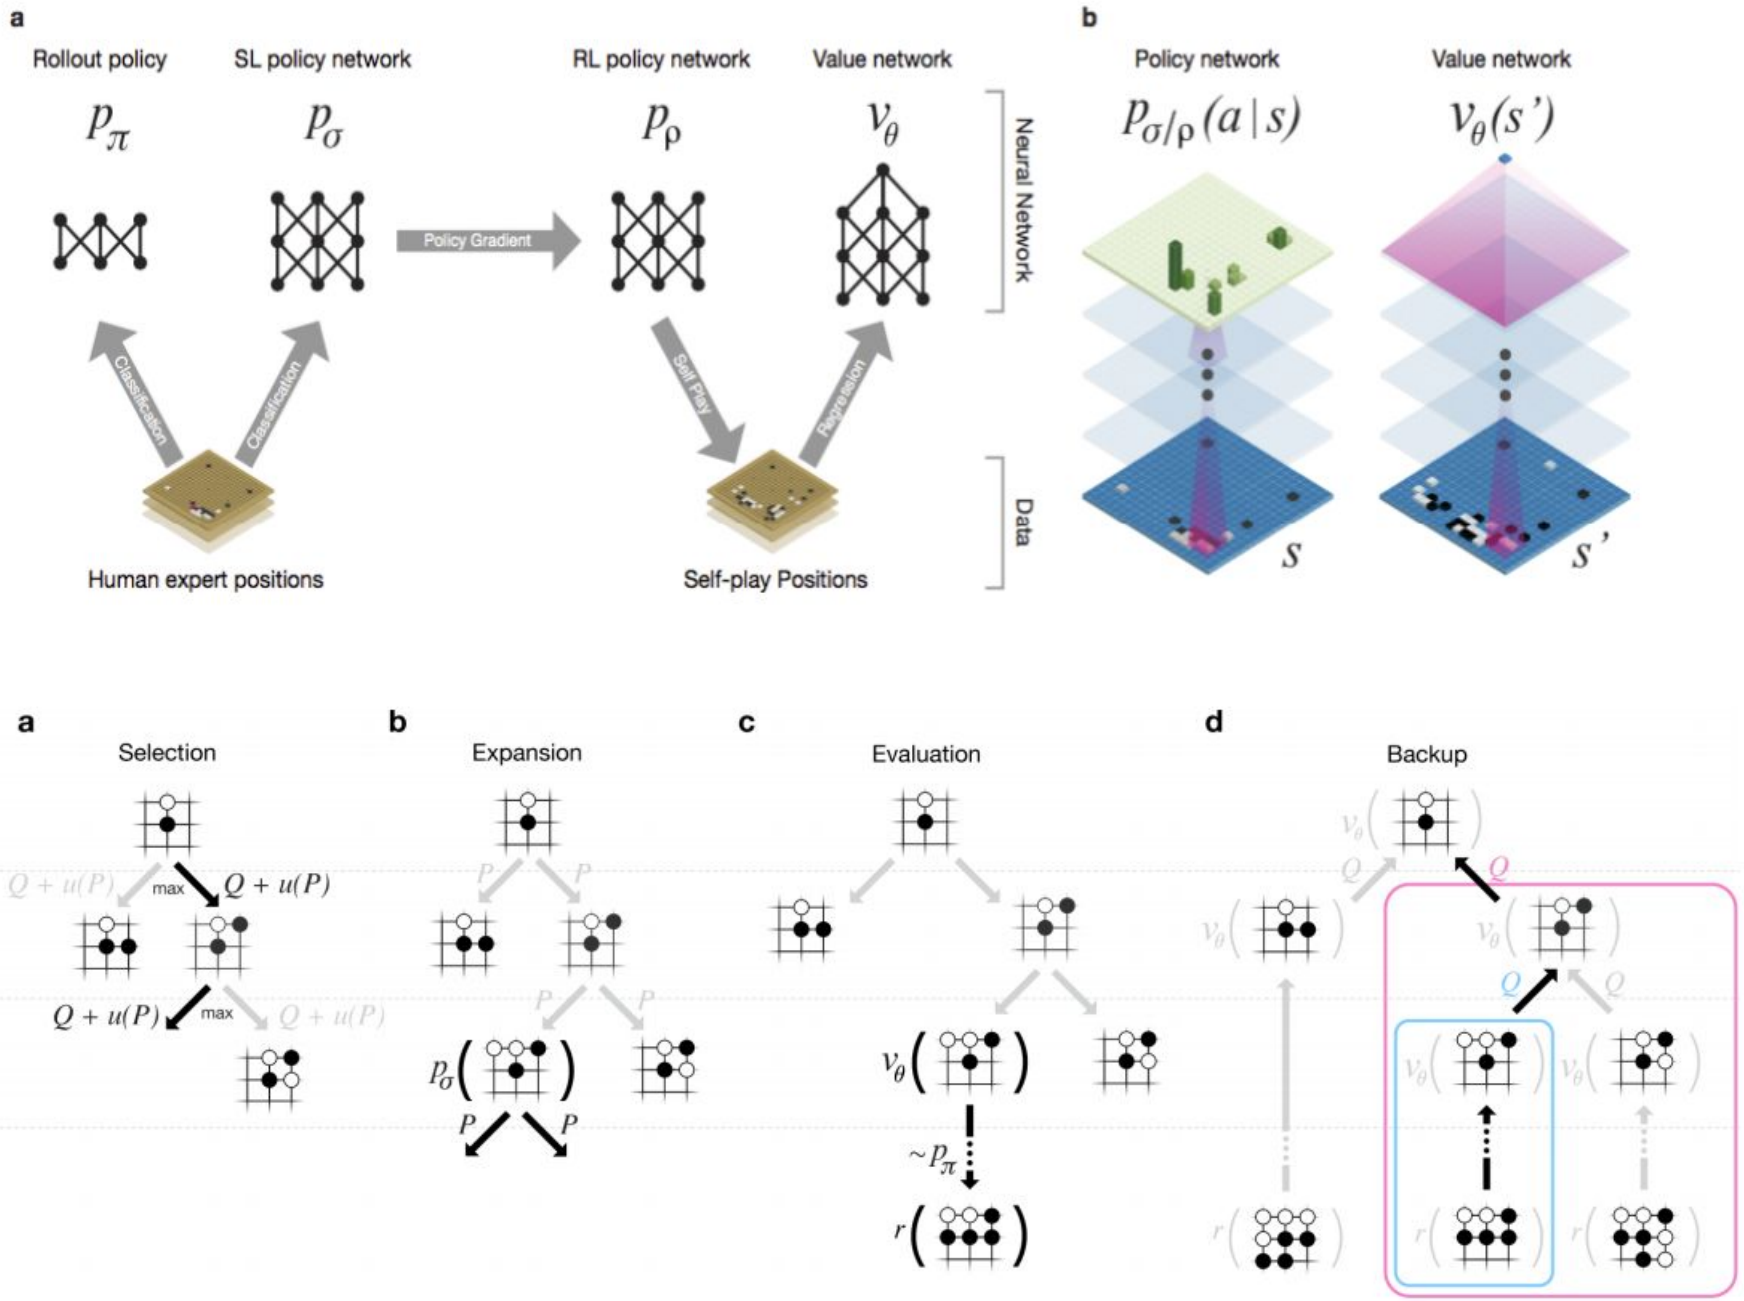
a
Rollout policy
SL policy network
RL policy network
Value network
p_π
p_σ
p_ρ
v_θ
Policy Gradient
Neural Network
Classification
Classification
Self Play
Regression
Data
Human expert positions
Self-play Positions
b
Policy network
Value network
$p_{\sigma/\rho}(a|s)$
$v_\theta(s')$
s
s'
a
Selection
$Q + u(P)$
max
$Q + u(P)$
$Q + u(P)$
max
$Q + u(P)$
b
Expansion
P
P
P
P
p_σ
P
P
c
Evaluation
v_θ
$\sim p_\pi$
r
d
Backup
v_θ
Q
Q
Q
Q
v_θ
v_θ
v_θ
v_θ
r
r
r

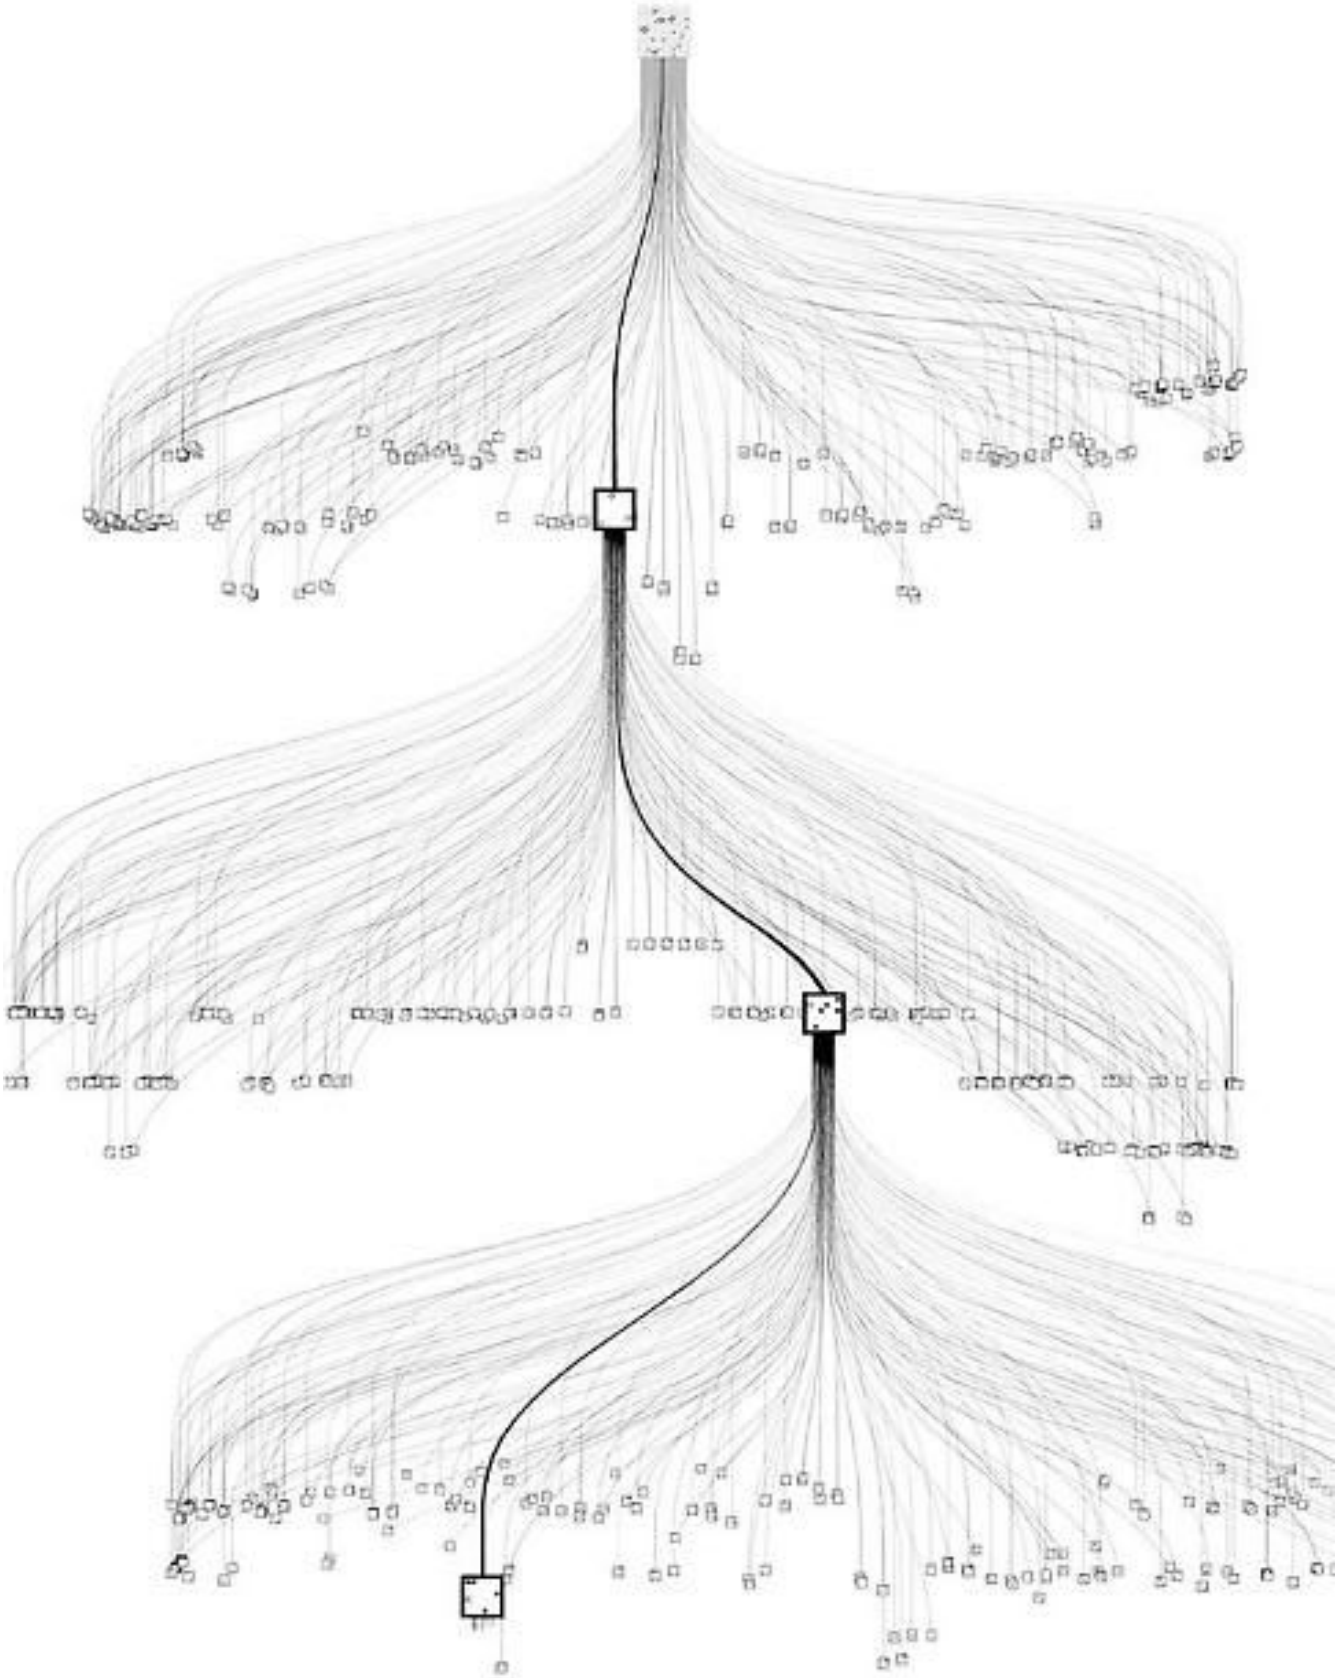

Компьютеры vs Человеки: Go

- AlphaGo
  - 1920 CPUs, 280 GPUs
  - 1 MW
  - ≈4 недели, 160,000 записей игр, 3,000,000 позиций
- Lee Sedol
  - 1 мозг (+ команда)
  - 100 W
  - 25+ лет карьеры

Тест Тьюринга, будущее

- If no one told me, maybe I would think the player was a little strange, but a very strong player, a real person.
- Планы по обучению сети «с нуля»

Центр управления полетами

Ресурсы

- https://habrahabr.ru/post/279071/
- http://neuralnetworksanddeeplearning.com/chap6.html
- http://www.cs.toronto.edu/~fritz/absps/imagenet.pd
- http://www.slideshare.net/DataScienceMD/deep-learning-with-gpus
- https://gogameguru.com/i/2016/03/deepmind-mastering-go.pdf
- http://jacquesmattheij.com/another-way-of-looking-at-lee-sedol-vs-alphago