Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CHAI
Humanities-Centered AI

Ralf Möller, Sylvia Melzer

AlphaGo, AlphaZero, AlphaGeometry, AlphaFold

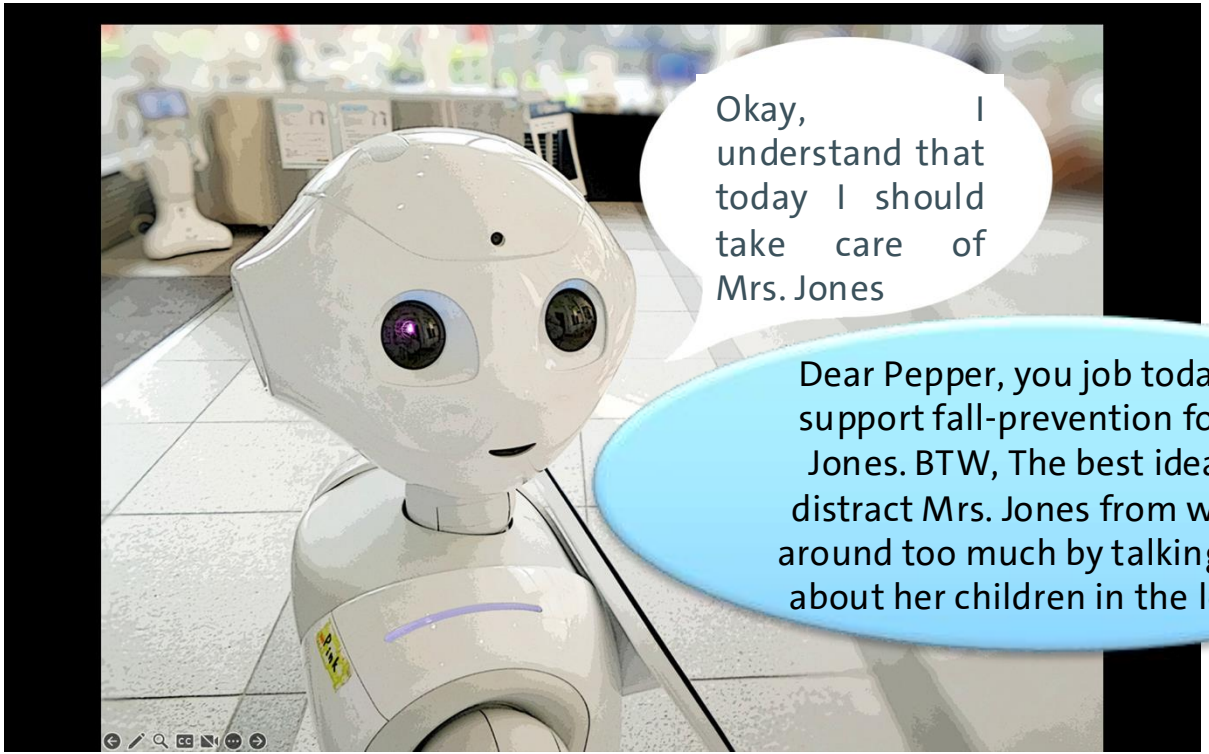
Okay, I understand that today I should take care of Mrs. Jones
Dear Pepper, you job today is to support fall-prevention for Mrs. Jones. BTW, The best idea is to distract Mrs. Jones from walking around too much by talking to her about her children in the lounge.

Intelligent Agent

Percepts

From percepts to the state of the environment (incl. task description)

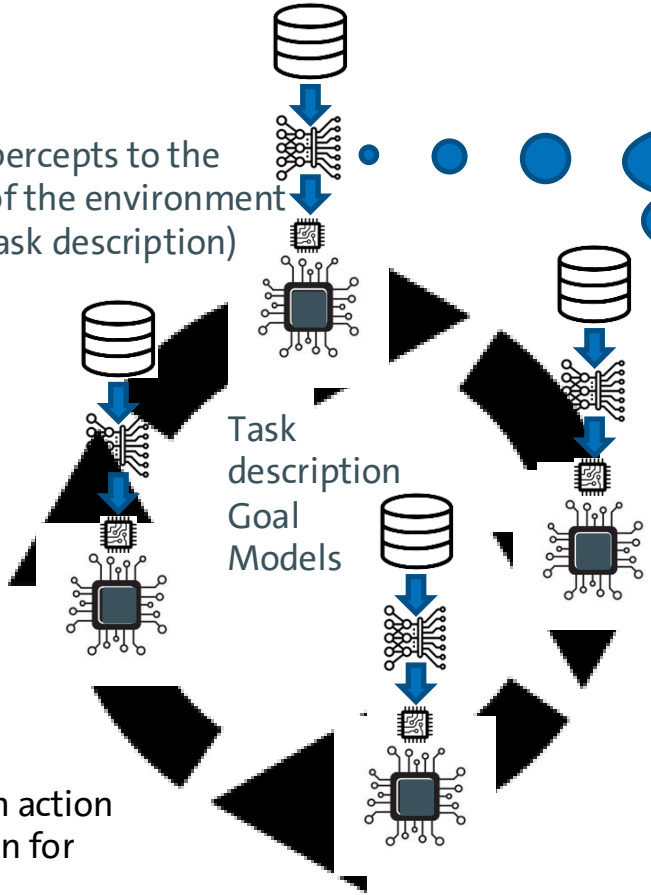

Intelligent agents should support multiple tasks (and ChatGPT does it)

Are the current goals still valid for task descriptions? Derive new goals if required.

Planning vs. Routine

What is the best action to reach the goals (carry out tasks)

Actions

Carry out or return action and/or explanation for human principal

Graceful degradation?

- AlphaGo – First program to defeat a professional Go player
- AlphaZero – More board games, less human feedback
- MuZero – Now playing Atari games, with not even the rules

https://youtu.be/IVMgxtm5L-U

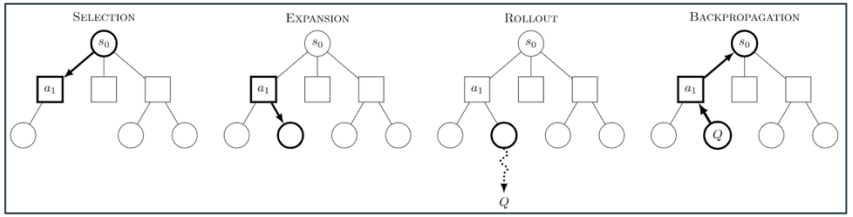

Monte Carlo Tree Search (MCTS)

- **Selection:** Start from root R and select successive child nodes until a leaf node L is reached. The root is the current game state and a leaf is any node that has a potential child from which no simulation (playout) has yet been initiated. The section below says more about a way of biasing choice of child nodes that lets the game tree expand towards the most promising moves, which is the essence of Monte Carlo tree search.
- **Expansion:** Unless L ends the game decisively (e.g. win/loss/draw) for either player, create one (or more) child nodes and choose node C from one of them. Child nodes are any valid moves from the game position defined by L.
- **Simulation:** Complete one random playout from node C. This step is sometimes also called playout or rollout. A playout may be as simple as choosing uniform random moves until the game is decided (for example in chess, the game is won, lost, or drawn).
- **Backpropagation:** Use the result of the playout to update information in the nodes on the path from C to R.

https://en.wikipedia.org/wiki/Monte_Carlo_tree_search

Monte Carlo Tree Search (MCTS)

- Example:

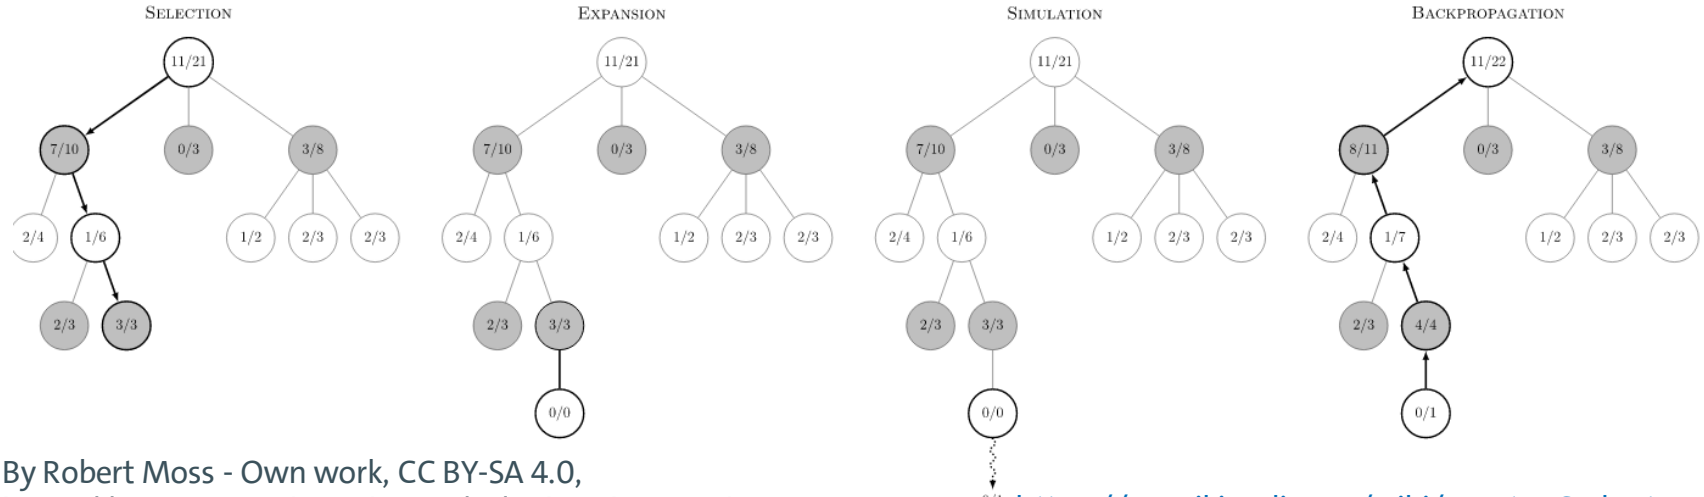

By Robert Moss - Own work, CC BY-SA 4.0,
https://commons.wikimedia.org/w/index.php?curid=88889583

https://en.wikipedia.org/wiki/Monte_Carlo_tree_search

Go

Introduction to Go

https://online-go.com/learn-to-play-go/rules-intro

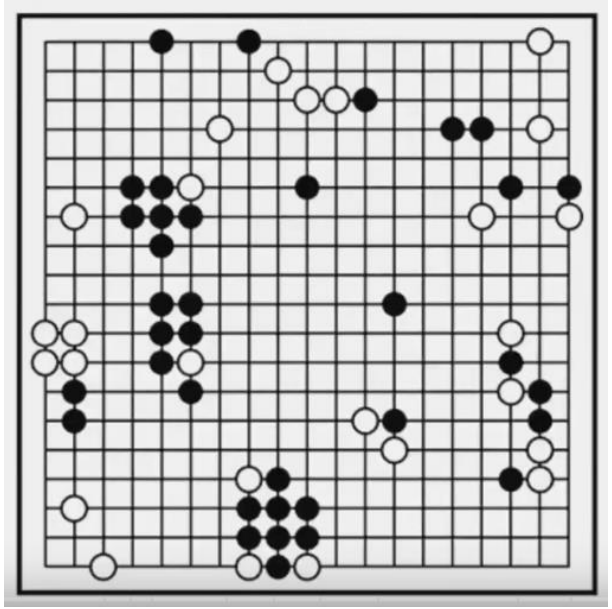

Go

- 19x19 lines and 361 intersections
- Very high game tree complexity

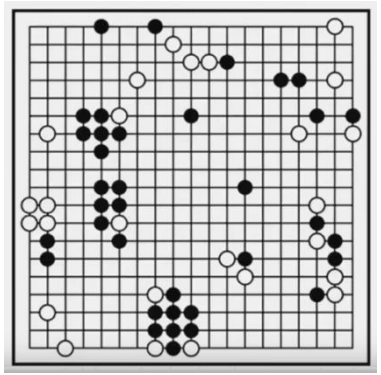

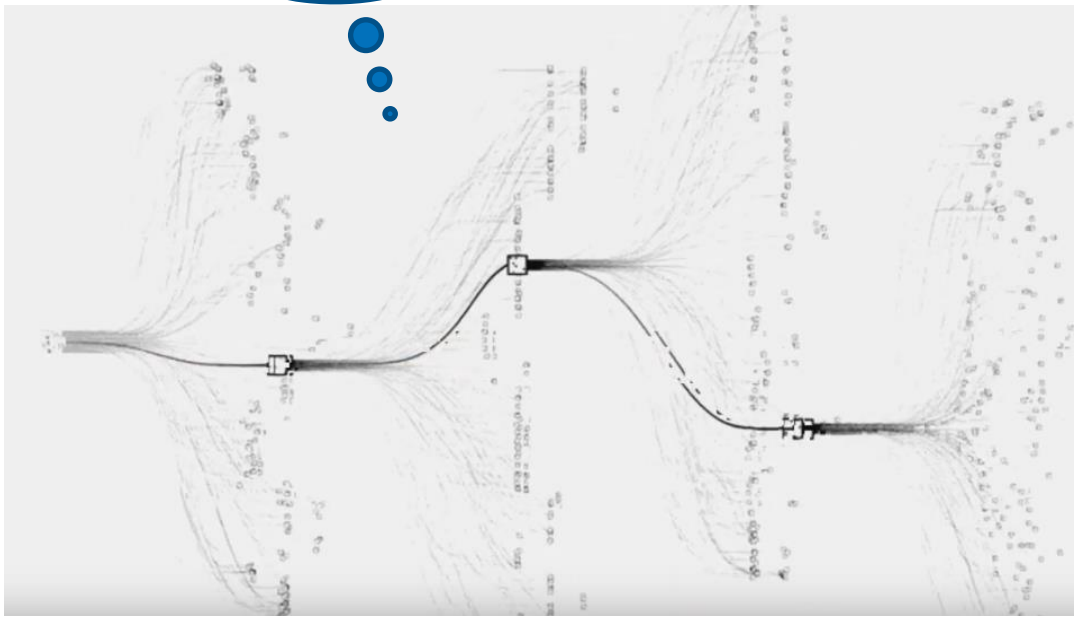

AlphaGo

- Cut down the complexity of the search tree via network policy

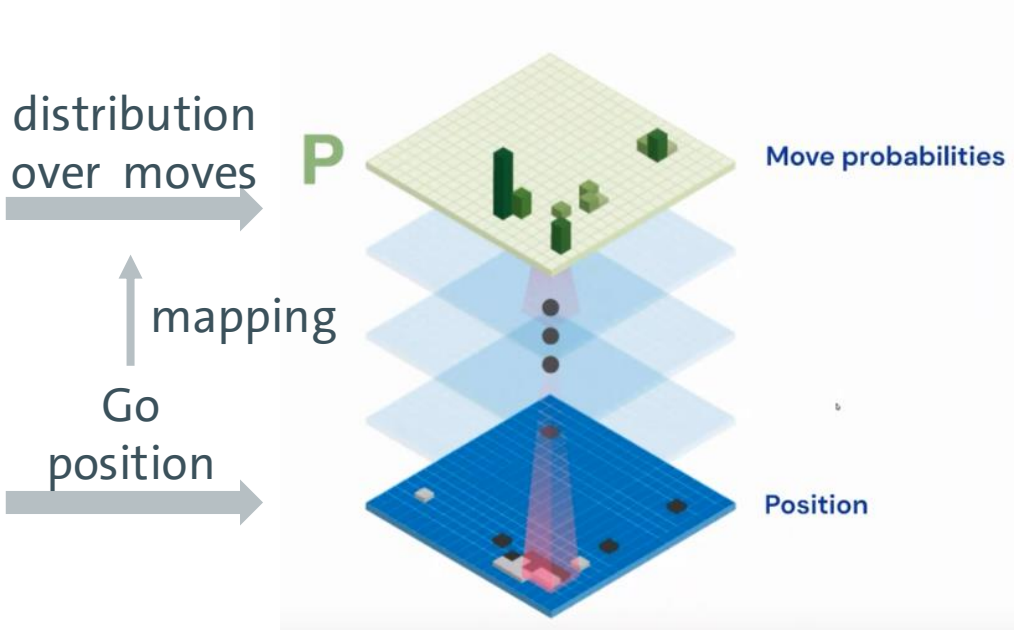

AlphaGo

- Evaluation network
- Feature vector represents the position
- Output: win/loss signal
- Here: black will win

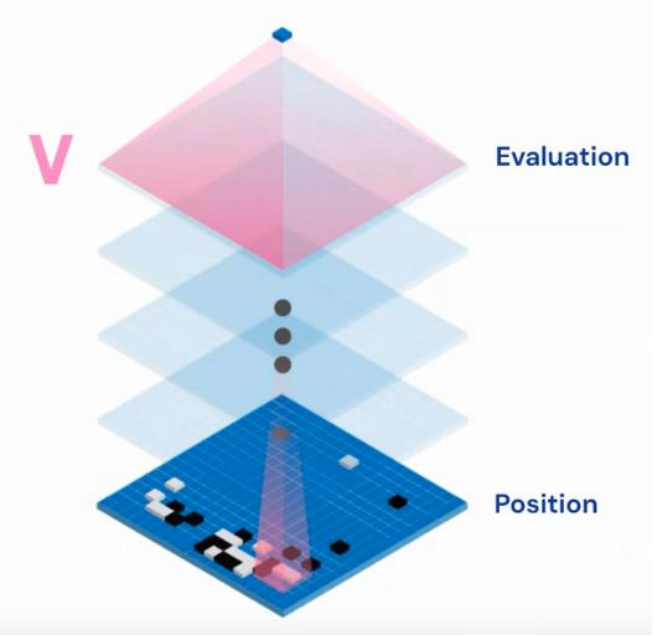

AlphaGo

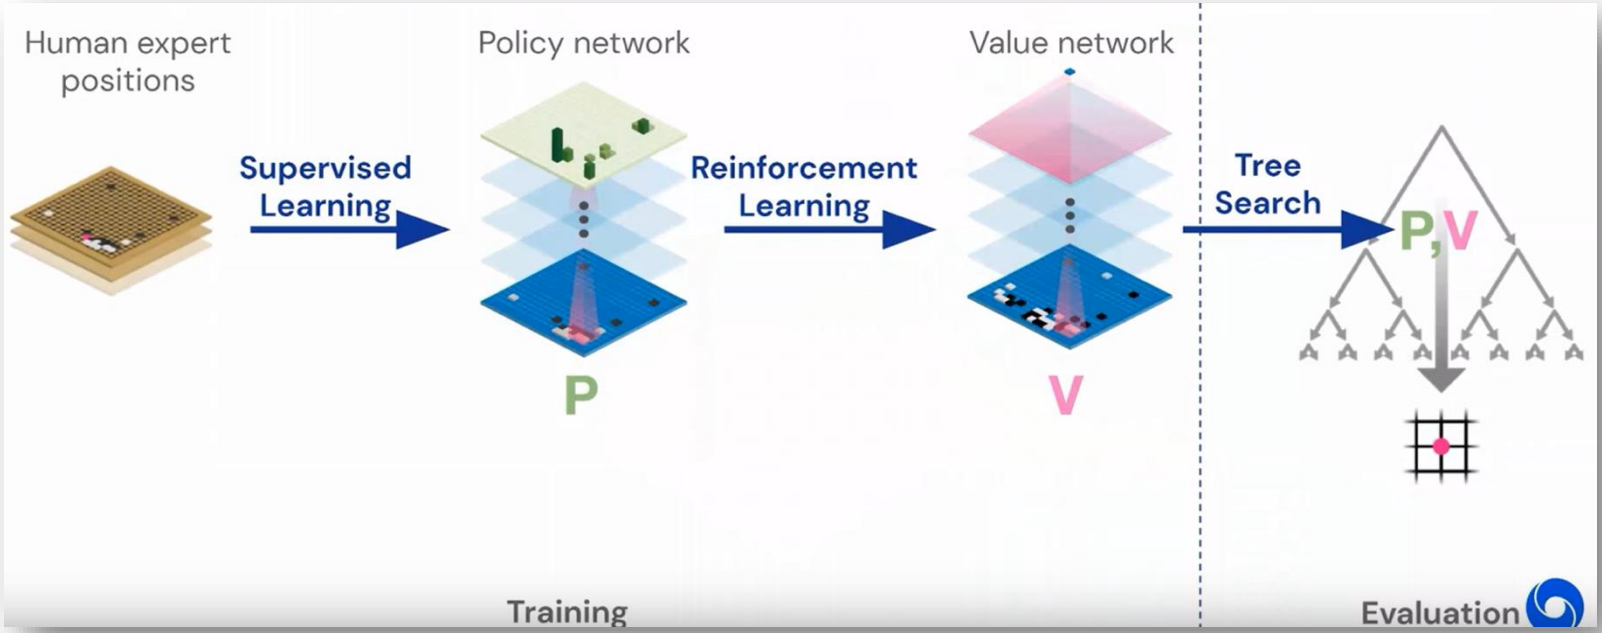
Human expert positions
Supervised Learning
Policy network
P
Reinforcement Learning
Value network
V
Tree Search
P,V
Training
Evaluation

AlphaGo

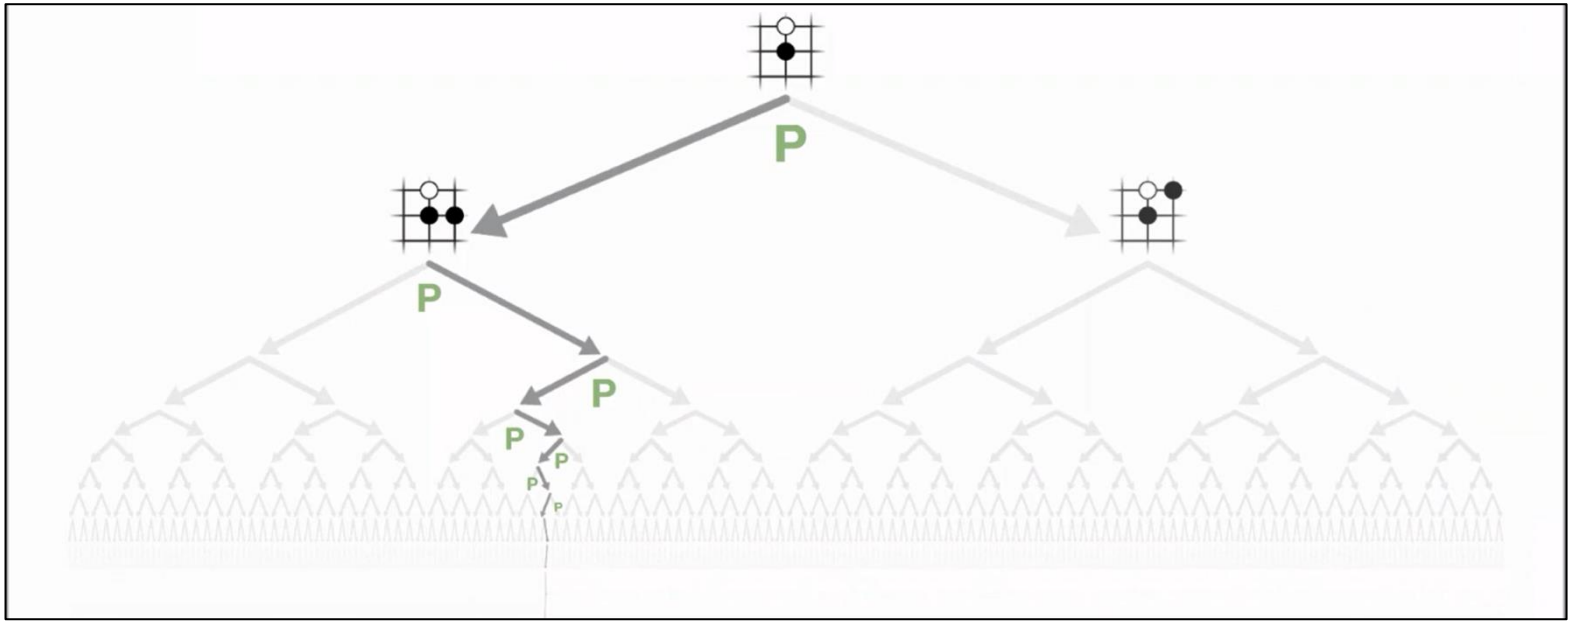

AlphaGo

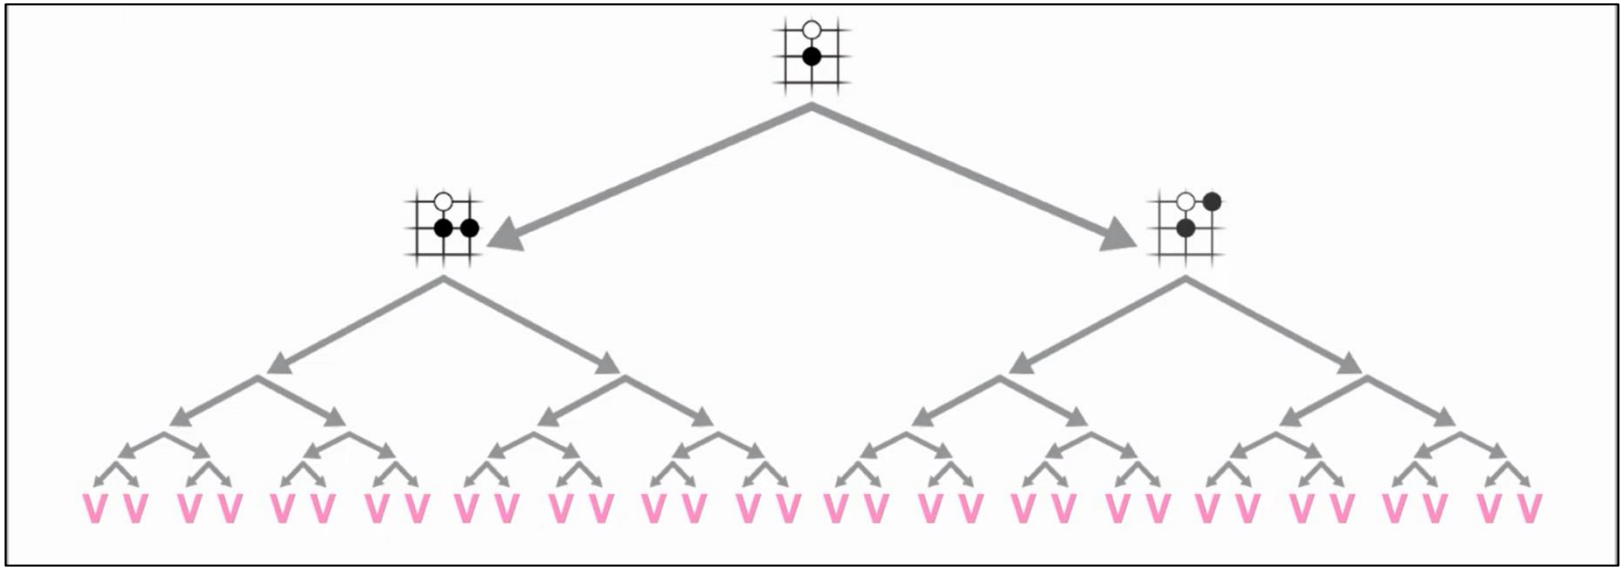

AlphaZero: One Algorithm, Three Games

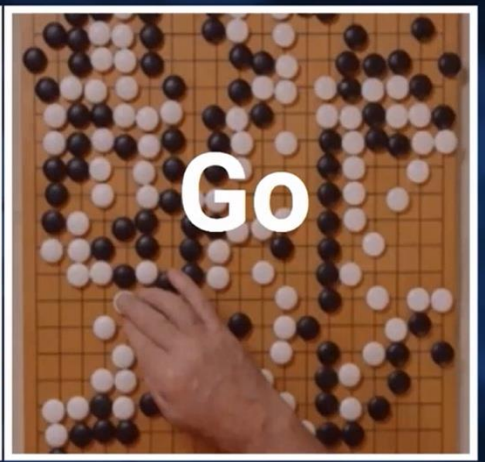

Silver, D., Schrittwieser, J., Simonyan, K. et al. Mastering the game of Go without human knowledge. Nature 550, 354–359 (2017).

Silver, D., Hubert, T., Schrittwieser, J. et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. Science 362(6419), pp. 1140–1144 (2018).

Graceful degradation?

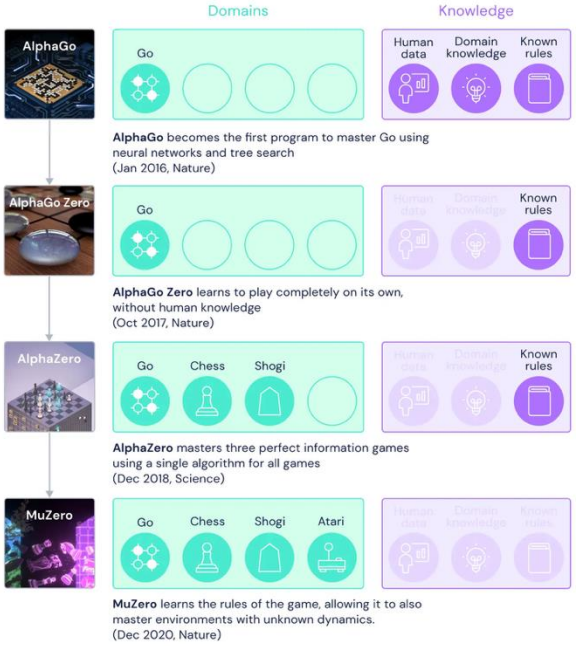

Monte Carlo Tree Search (MCTS)

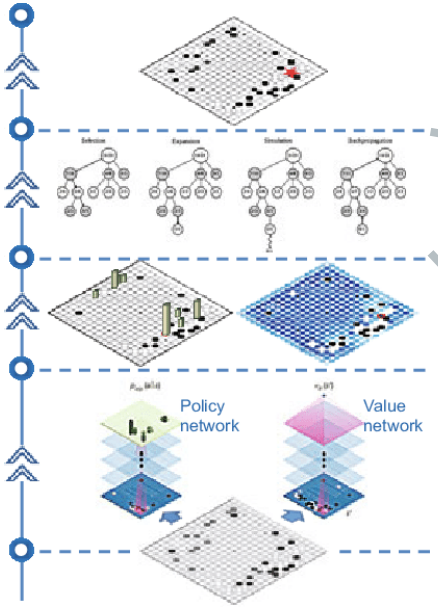

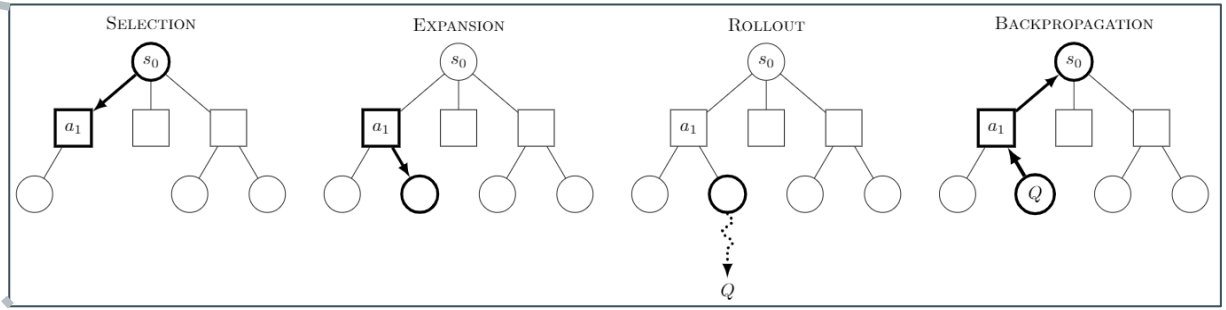

https://en.wikipedia.org/wiki/Monte_Carlo_tree_search

Agains the Brittleness: AlphaZero

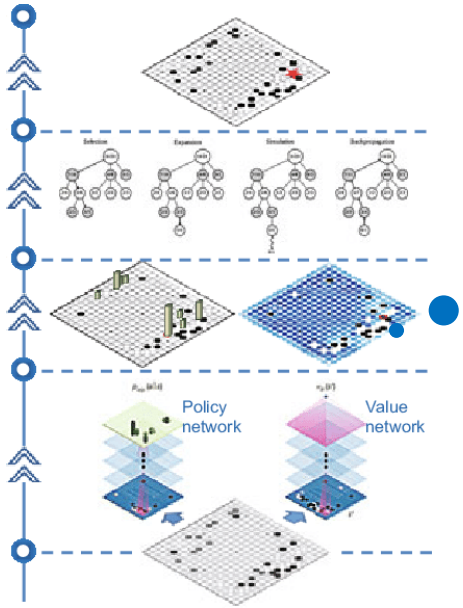

What are good starting points for policies?

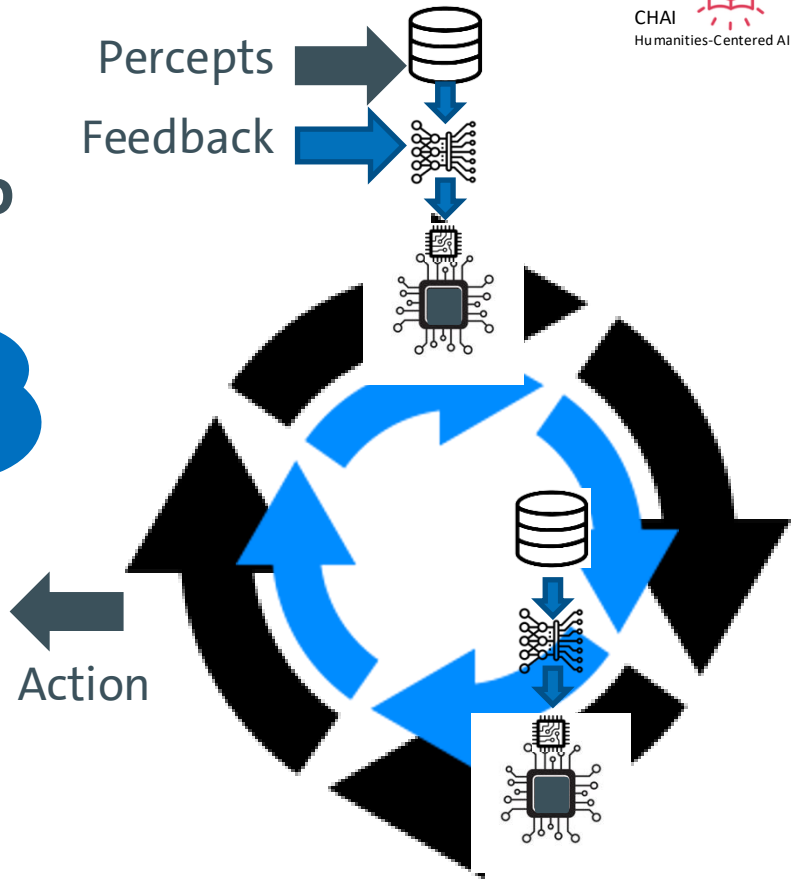

AlphaZero plays games against itself

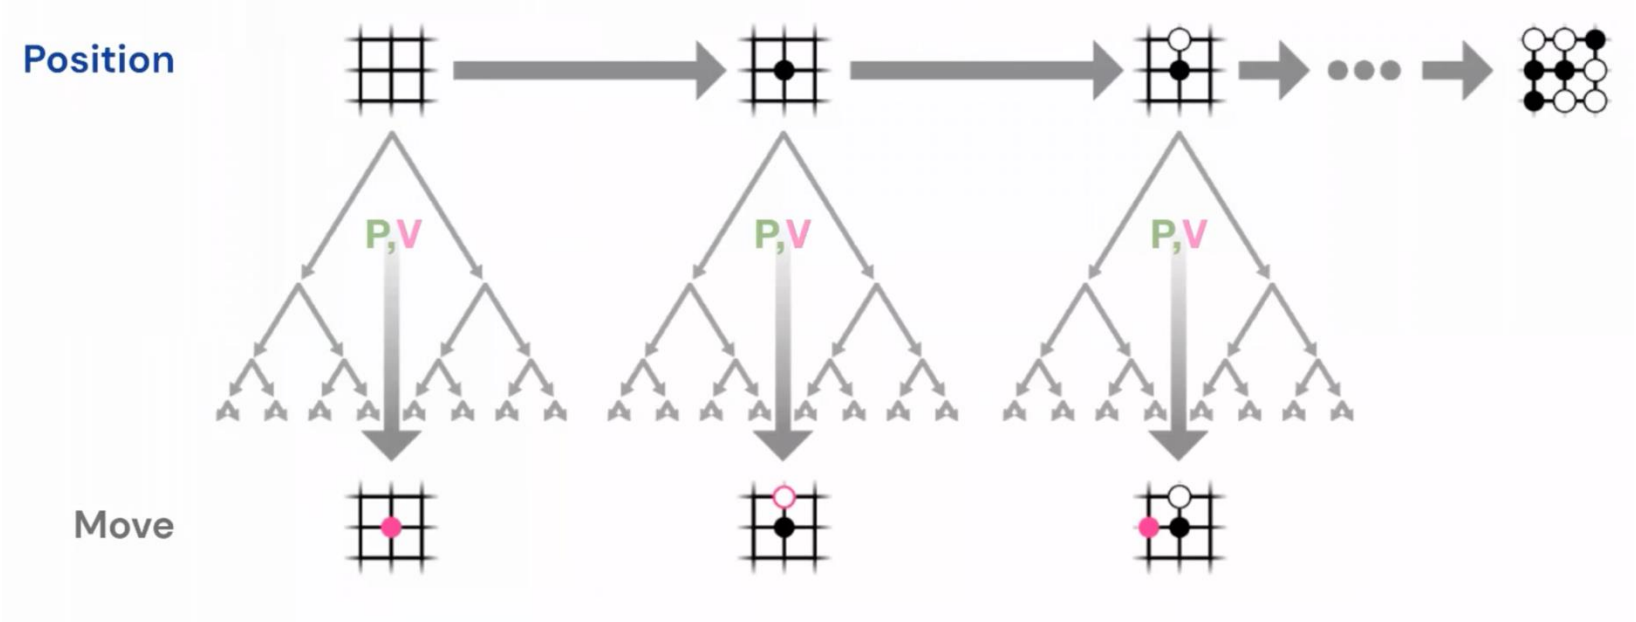

Reinforcement Learning in AlphaZero

New value network V’ is trained to predict winner

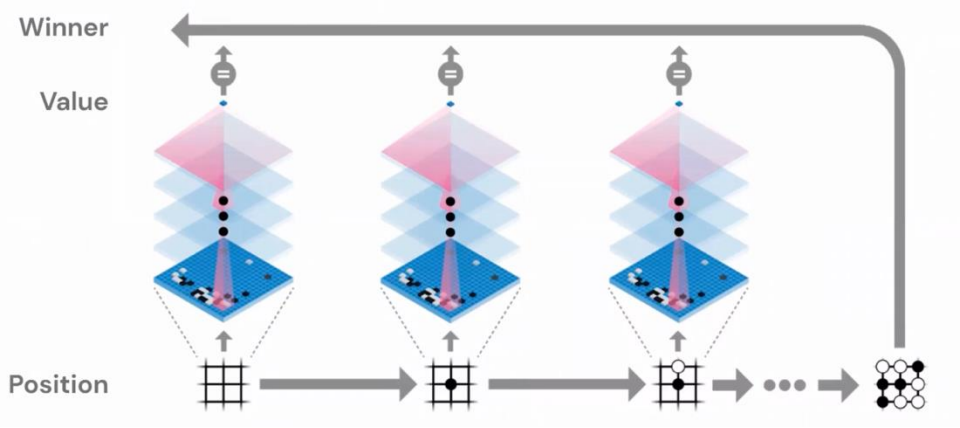

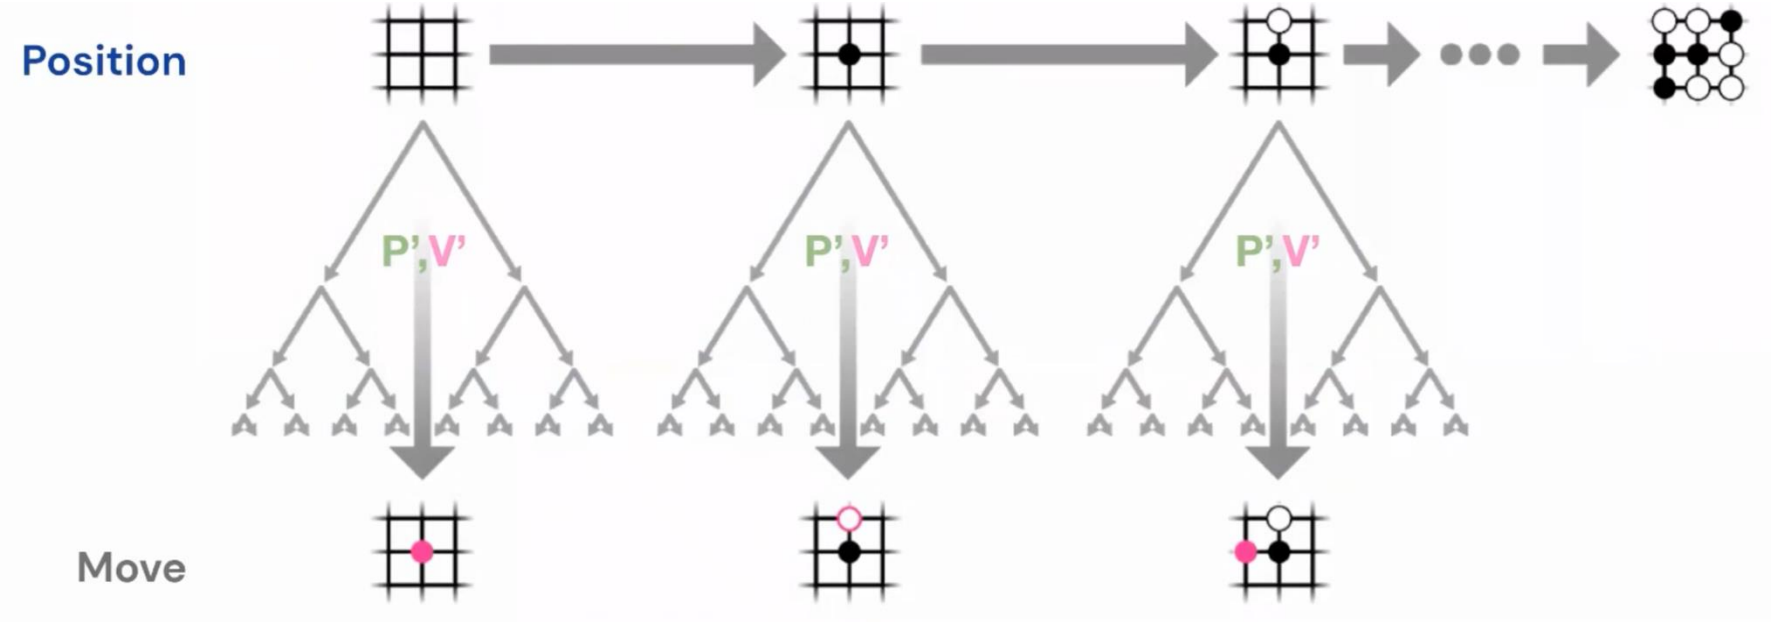
Position
P',V'
P',V'
P',V'
Move

Representation Learning

- What is a good representation for the environment given a set of task descriptions
- There might be representations suggested from task descriptions
- There might be things to be added for supporting specific problem solving strategies effectively

Find representation augmentations: AlphaGeometry

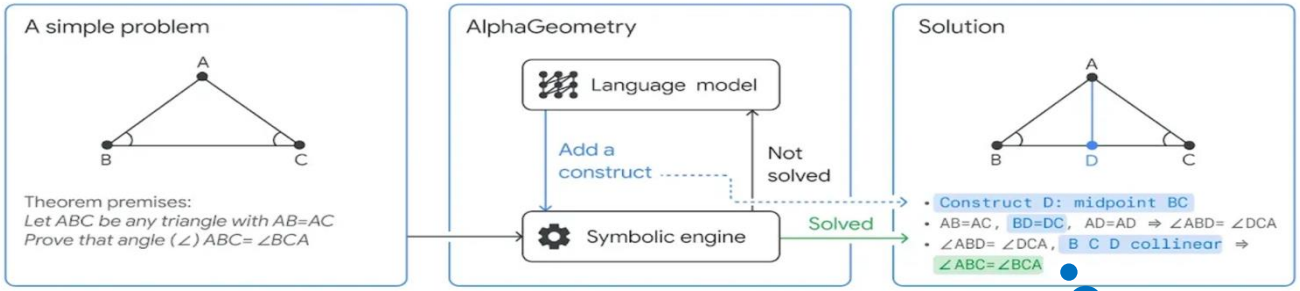

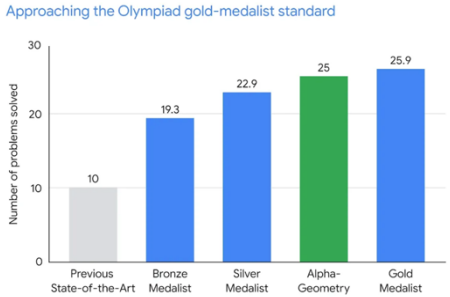

No chain of thought (CoT) reasoning
Need to deal with case distinctions!

9 dots puzzle: Creative solutions

Connect with four (straight) lines w/o releasing pen

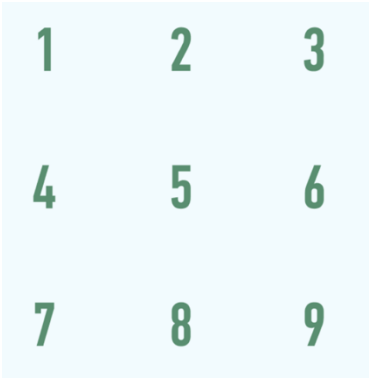

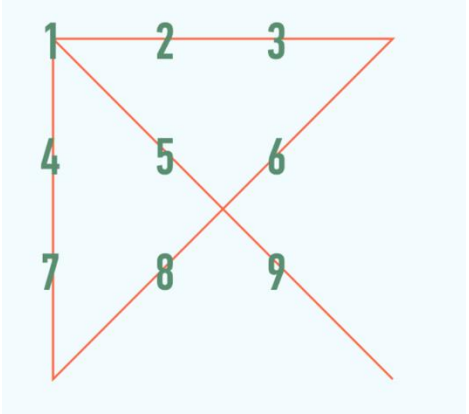

Start thinking “out of the box”

LLMs as providers for ideas that are rigorously checked

Prediction is key (see Friston et al.)

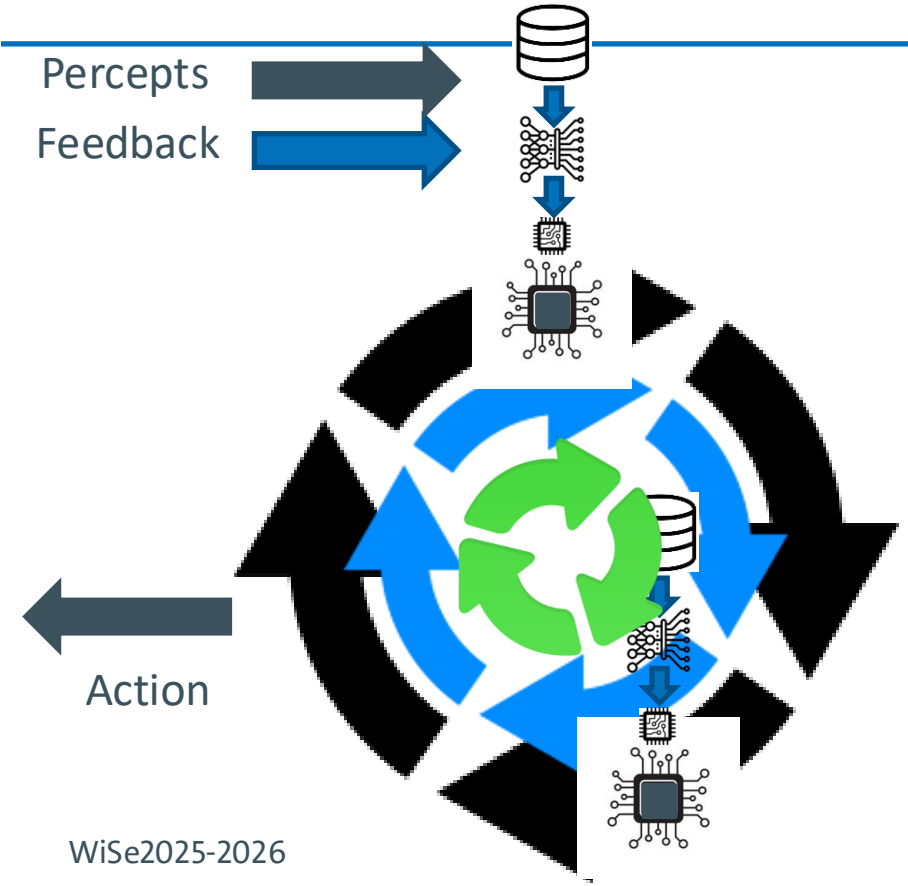

- What about executing a multistep plan?
- Execute respective next best action
  - Could be an action from a multistep plan
- Observe environment
  - ... which might change stochastically (and due to other agents)
- Prediction
  - It might become clear for an agent that subsequent actions from specific (incomplete) plans
  - are never applicable or
  - never contribute to goals
- Purge pending actions
- Perplexity Minimization
  - Guide perception

A humanoid robot might be used to carry out different tasks

However, usually it is a bad idea to have robots with far from optimal physical design for a specific tasks

Powerful Industrial Robotics

Alphafold as the paradigm of a New Era

- Proteins are built from 22 standard amino acids
- HP35 is one of the smallest known naturally occurring protein (35 amino acids long) that folds into a defined, stable three-dimensional structure (a three-helix bundle) on its own
- The number of possible sequences for a 35 amino-acid chain is: 20^35
- 38,629,697,920,000,000,000,000,000,000,000,000,000,000,000,000
- Median protein length of human: 416
- Problem: Find out structures is slow and expensive

https://doi.org/10.1063/1.2995987,
https://bionumbers.hms.harvard.edu/files/Median%20protein%20lengths%20in%20eukaryotic,%20bacterial%20and%20archaeal%20organisms.pdf

AlphaFold

The number of protein structures known (post-Alpha Fold 2): 200,000,000

Our malaria vaccine work highlighted by AlphaFold

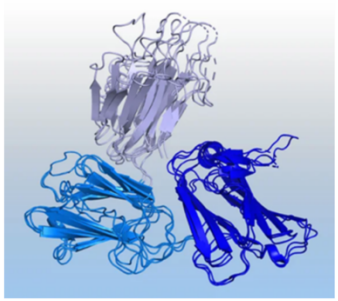

"Matt Higgins and his team of researchers at the University of Oxford had a problem", say Time magazine, and AlphaFold2 helped us to solve it.

It was a pleasure to appear at the press conference in which AlphaFold2 announced a massive database containing predictions of protein structures for all sequenced organisms. We were there to describe how AlphaFold2, combined with experimental data, allowed us to understand the malaria transmission-blocking antibody candidate Pfs48/45.

July 2022

This work is described in our preprint:

Ko, K.T., Lennartz, F., Mekhaiel, D., Guloglu, B., Marini, A., Deuker, D.J., Long, C.A., Jore, M.M., Miura, M., Biswas, S. and Higgins, M.K. (2022) Structure of the malaria vaccine candidate Pfs48/45 and its recognition by transmission blocking antibodies. Nature Communications 13 5603

Bertoline LMF, Lima AN, Krieger JE, Teixeira SK. Before and after AlphaFold2: An overview of protein structure prediction. Front Bioinform. **2023** Feb 28;3:1120370. doi: 10.3389/fbinf.2023.1120370. PMID: 36926275; PMCID: PMC10011655.

https://higginslab.web.ox.ac.uk/our-malaria-vaccine-work-highlighted-alphafold

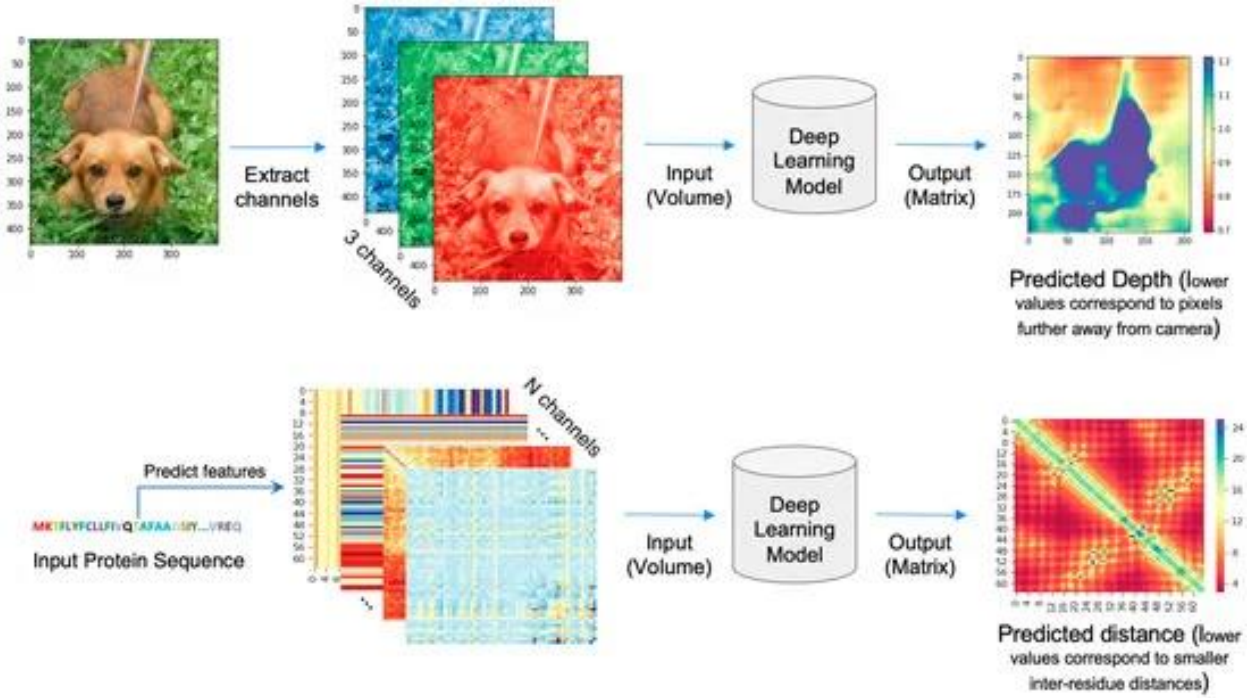

Pakhrin, Subash C., Bikash Shrestha, Badri Adhikari, and Dukka B. Kc. 2021. "Deep Learning-Based Advances in Protein Structure Prediction." International Journal of Molecular Sciences 22 (11): 5553. https://doi.org/10.3390/ijms22115553

AlphaFold 2

One sequence

Multiple Sequence Alignment (MSA)

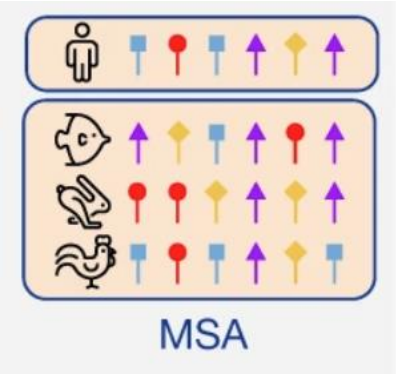

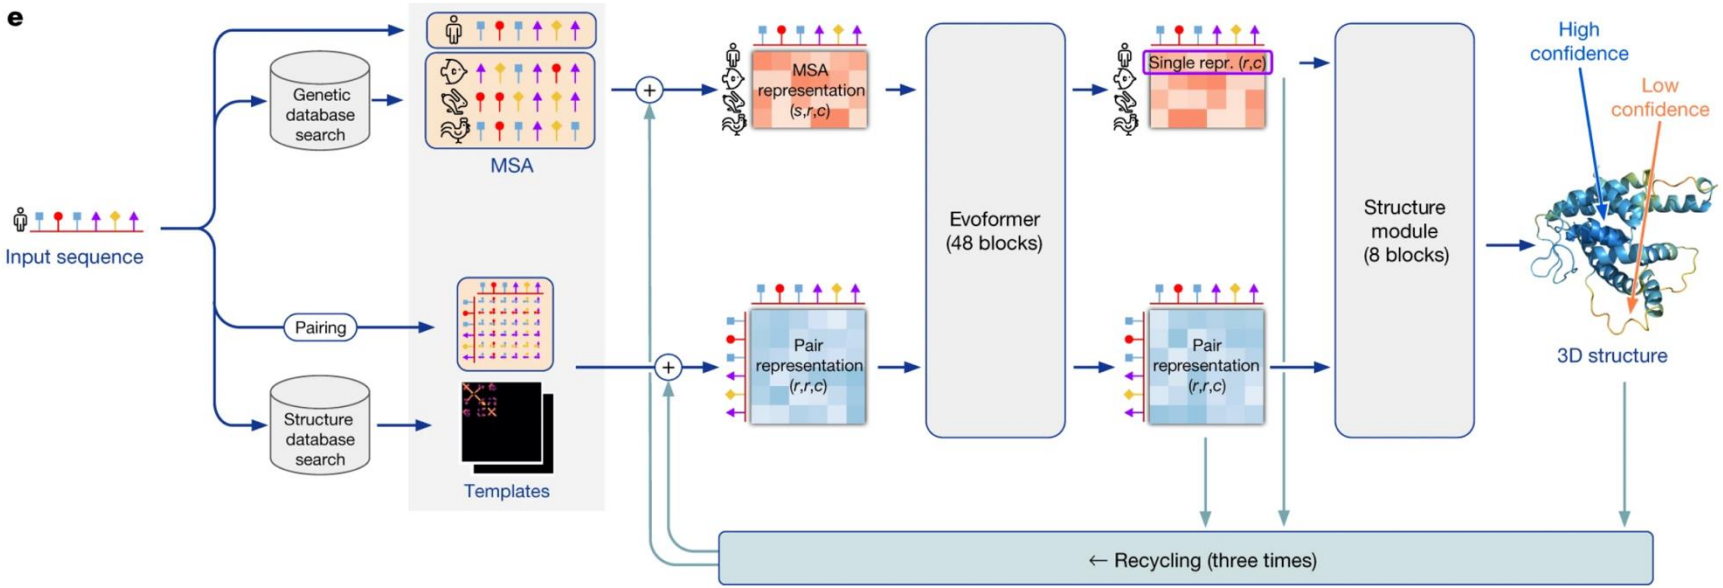

Jumper, J., Evans, R., Pritzel, A. *et al*. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021). https://doi.org/10.1038/s41586-021-03819-2

AlphaFold

Demo: https://alphafold.com/

- What does AlphaFold do well?
- What are its limitations?

AlphaFold

Strengths:

- Provides fast access to predicted protein structures.
- High accuracy for many proteins.
- No experimental work is required by the user.

Limitations:

- Structures are **predictions**, not experimental measurements.
- Protein dynamics, ligand binding, and multi-protein complexes are not fully captured.
- Some regions have low prediction confidence.