


HOW CHESS ENGINES WORK




ANVAY SHAH


FOR DEEP RL WORKSHOP, SPRING 2024 (PROF. HARSHAD KHADILKAR)



HISTORY OF ENGINES (AND HUMANS)



HISTORY OF ENGINES (AND HUMANS)



CHESS ENGINES BEFORE NEURAL NETS

HAND-CRAFTED EVAL (HCE)

Question: (given a chess position) How do you know who is winning?

CHESS ENGINES BEFORE NEURAL NETS

HAND-CRAFTED EVAL (HCE)

Question: (given a chess position) How do you know who is winning?

Heuristic-based HCE functions - $V(s)$

Weighted sum of domain specific quantities - linear function approximation

CHESS ENGINES BEFORE NEURAL NETS

HAND-CRAFTED EVAL (HCE)

Question: (given a chess position) How do you know who is winning?

Heuristic-based HCE functions - $V(s)$

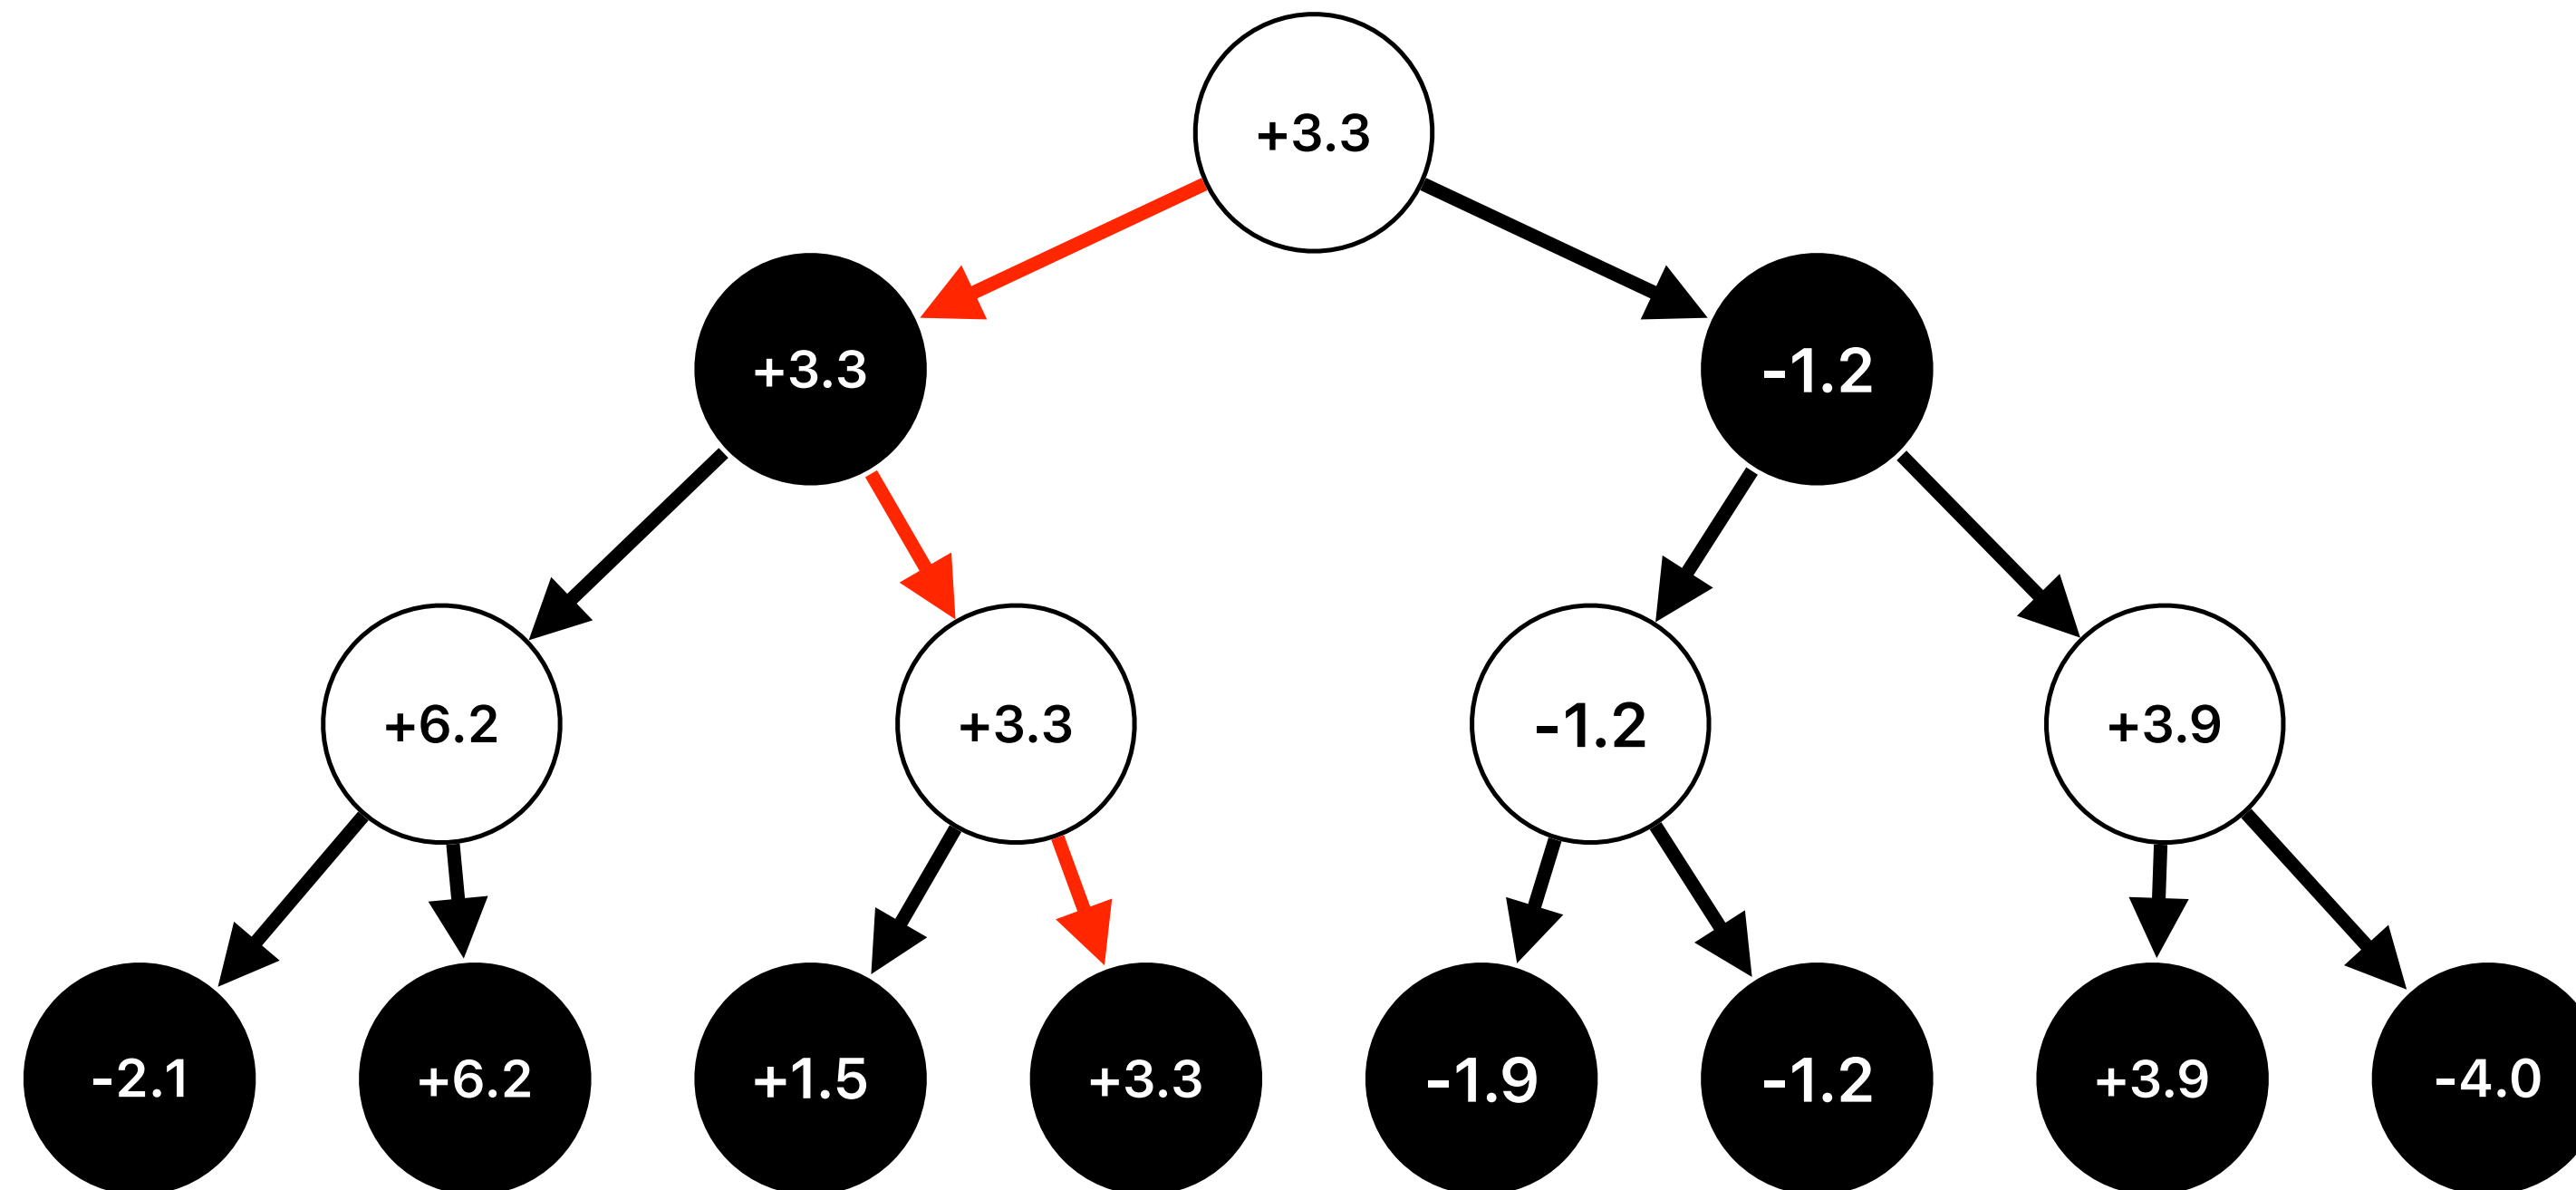
Weighted sum of domain specific quantities - linear function approximation

Material imbalance, positional advantage, material advantage, strategical advantage for pawns, strategical advantage for other pieces, incoming threats, passed pawns, space, king safety...

Tapered Eval: Different set of weights for different stages of the game

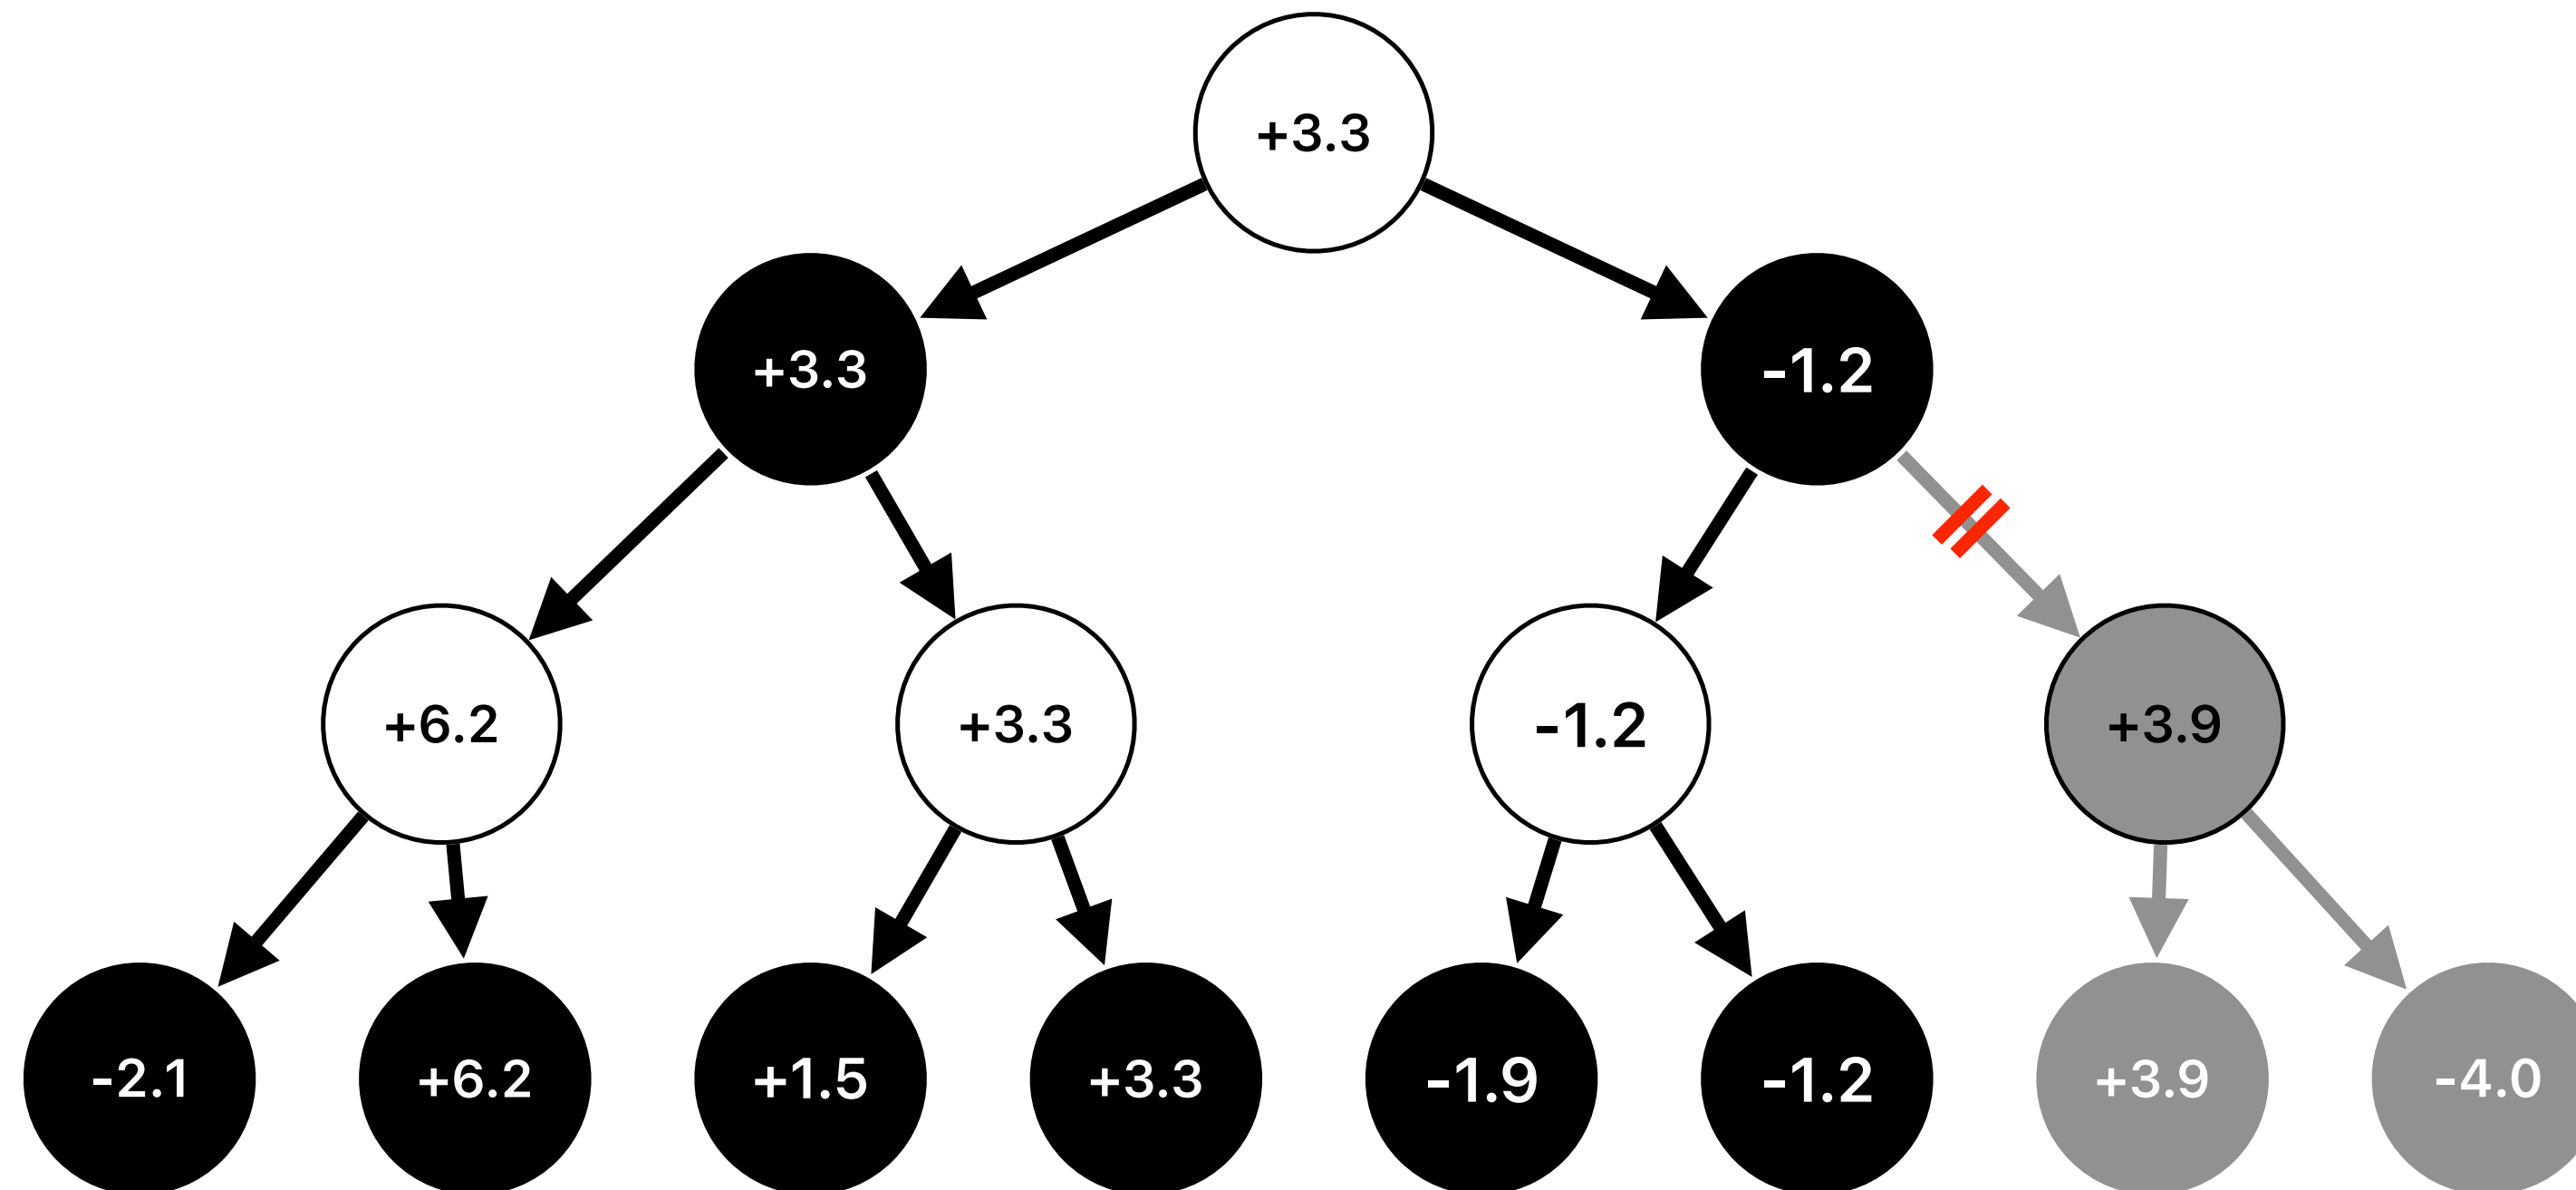
CHESS ENGINES BEFORE NEURAL NETS

MINIMAX SEARCH



CHESS ENGINES BEFORE NEURAL NETS

ALPHA-BETA PRUNING



ALPHAZERO GIVES STOCKFISH A BEATING

DEEP REINFORCEMENT LEARNING, SELF-PLAY

Learned to play Chess with zero domain knowledge through self-play

28 wins, 72 draws, 0 losses in 100 games vs Stockfish 8

Exhibited never seen before strategies and creative play in Chess

ALPHAZERO GIVES STOCKFISH A BEATING

DEEP REINFORCEMENT LEARNING, SELF-PLAY

Deep convolutional network with value and policy head - $V(s)$ and $\pi(s)$

Trained through self-play over 1000s of games - single net maintained

ALPHAZERO GIVES STOCKFISH A BEATING

DEEP REINFORCEMENT LEARNING, SELF-PLAY

Deep convolutional network with value and policy head - $V(s)$ and $\pi(s)$

Trained through self-play over 1000s of games - single net maintained

Question: How does AlphaZero choose a move?

ALPHAZERO GIVES STOCKFISH A BEATING

DEEP REINFORCEMENT LEARNING, SELF-PLAY

Deep convolutional network with value and policy head - $V(s)$ and $\pi(s)$

Trained through self-play over 1000s of games - single net maintained

Question: How does AlphaZero choose a move?

Monte Carlo Tree Search (MCTS) + Predictor Upper Confidence Tree Search (PUCT)

Choose action that maximises $Q(s, a)$ (exploitation) + U (exploration)

ALPHAZERO GIVES STOCKFISH A BEATING

WHERE IS ALPHAZERO TODAY?

Leela Chess Zero - AlphaZero framework with different model architectures

Top Chess Engine Championship (TCEC) - LCZero usually 2nd behind Stockfish in the last 5 years, occasionally 1st or 3rd

STOCKFISH 12 (AND ABOVE)

NNUE (EFFICIENTLY UPDATABLE NEURAL NETWORK)



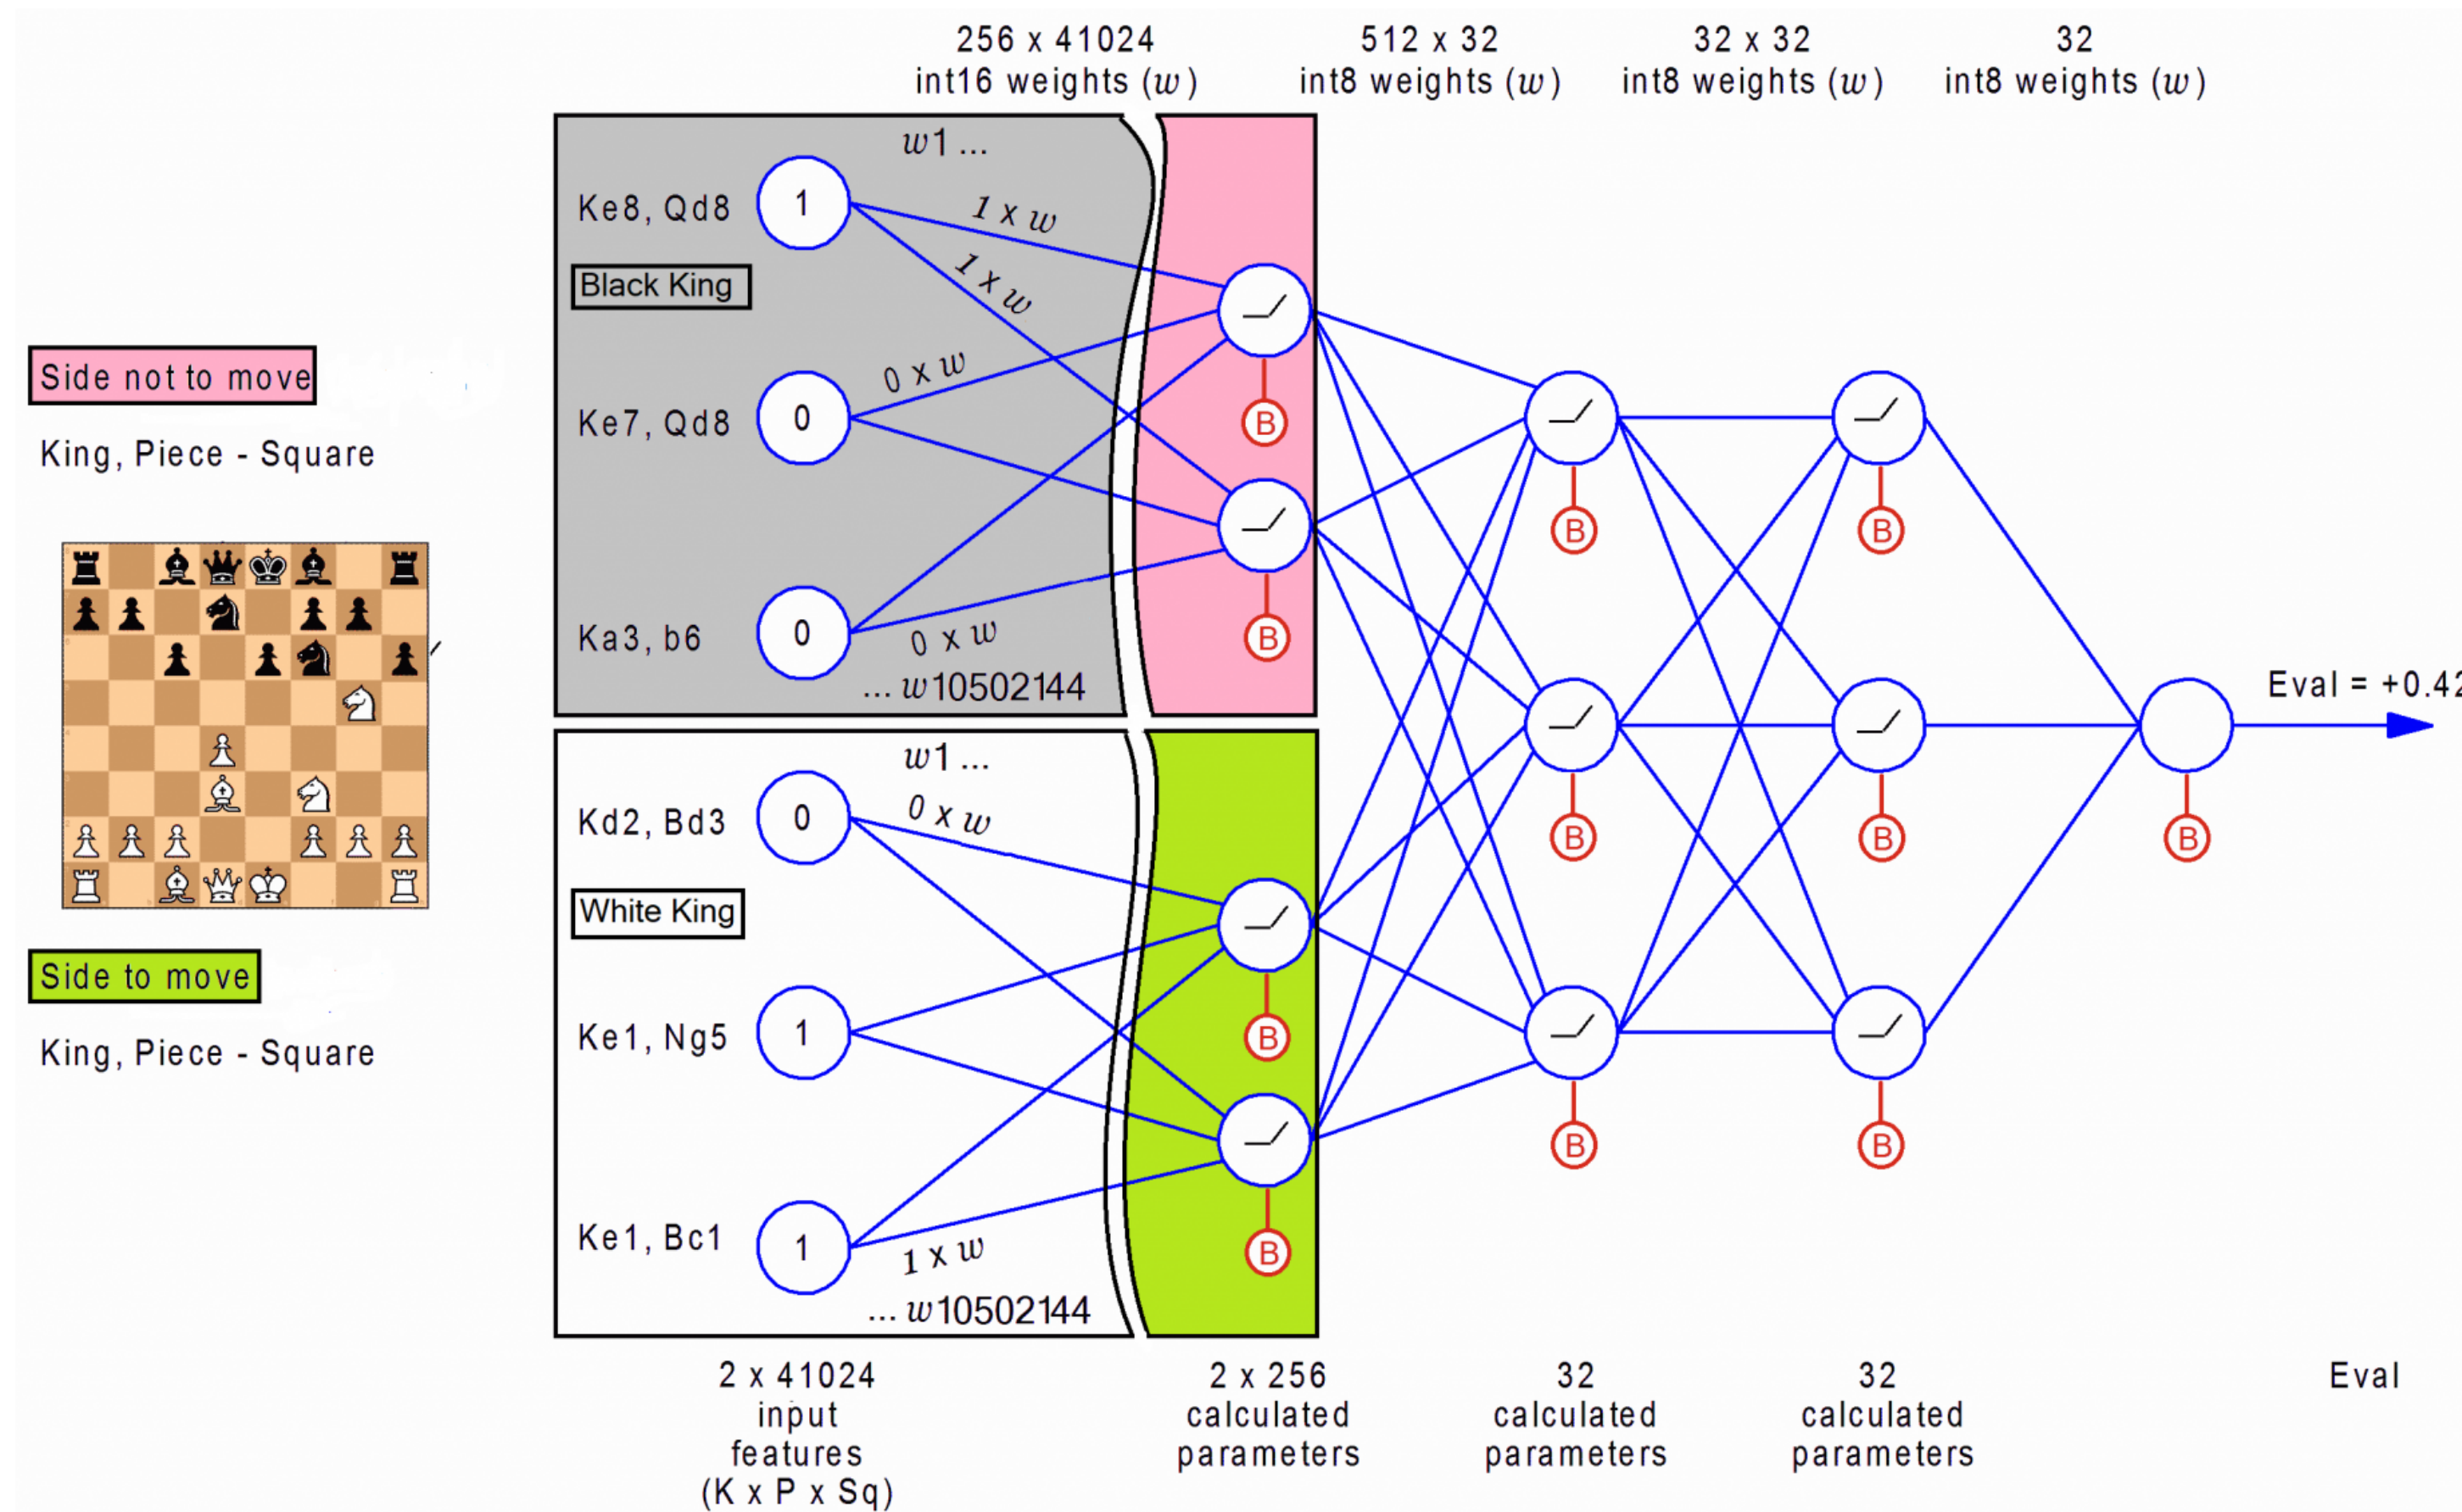
Idea: Use a small (4 layers) Neural Network for state evaluation in search

Earlier used successfully in Shogi, introduced in chess since Stockfish 12

Slower but **80 Elo stronger** than HCE

NNUE ARCHITECTURE


HALFKP STRUCTURE



Question: How to encode a chess position into features for a NN?

NNUE ARCHITECTURE

HALFKP STRUCTURE




Question: How to encode a chess position into features for a NN?

Naive way - A feature for each tuple - (square, piece, colour)

Number of features = $64 \times 6 \times 2 = 768$

NNUE ARCHITECTURE

HALFKP STRUCTURE




HalfKP input is overparametrized

(king, square, piece, colour) x 2

Number of features = $64 \times 64 \times 5 \times 2 \times 2$
= 40960 x 2

NNUE ARCHITECTURE

HALFKP STRUCTURE



Incremental updates: fast forward pass as you only need to perform matrix multiplications for the features which have changed

Training: Supervised learning on Stockfish 11 evaluations

STOCKFISH 12 (AND ABOVE)

NNUE (EFFICIENTLY UPDATABLE NEURAL NETWORK)



Stockfish 14 - HalfKAv2 architecture used instead of HalfKP

Stockfish 16 - HCE removed entirely and replaced with NNUE

LEELA CHESS ZERO



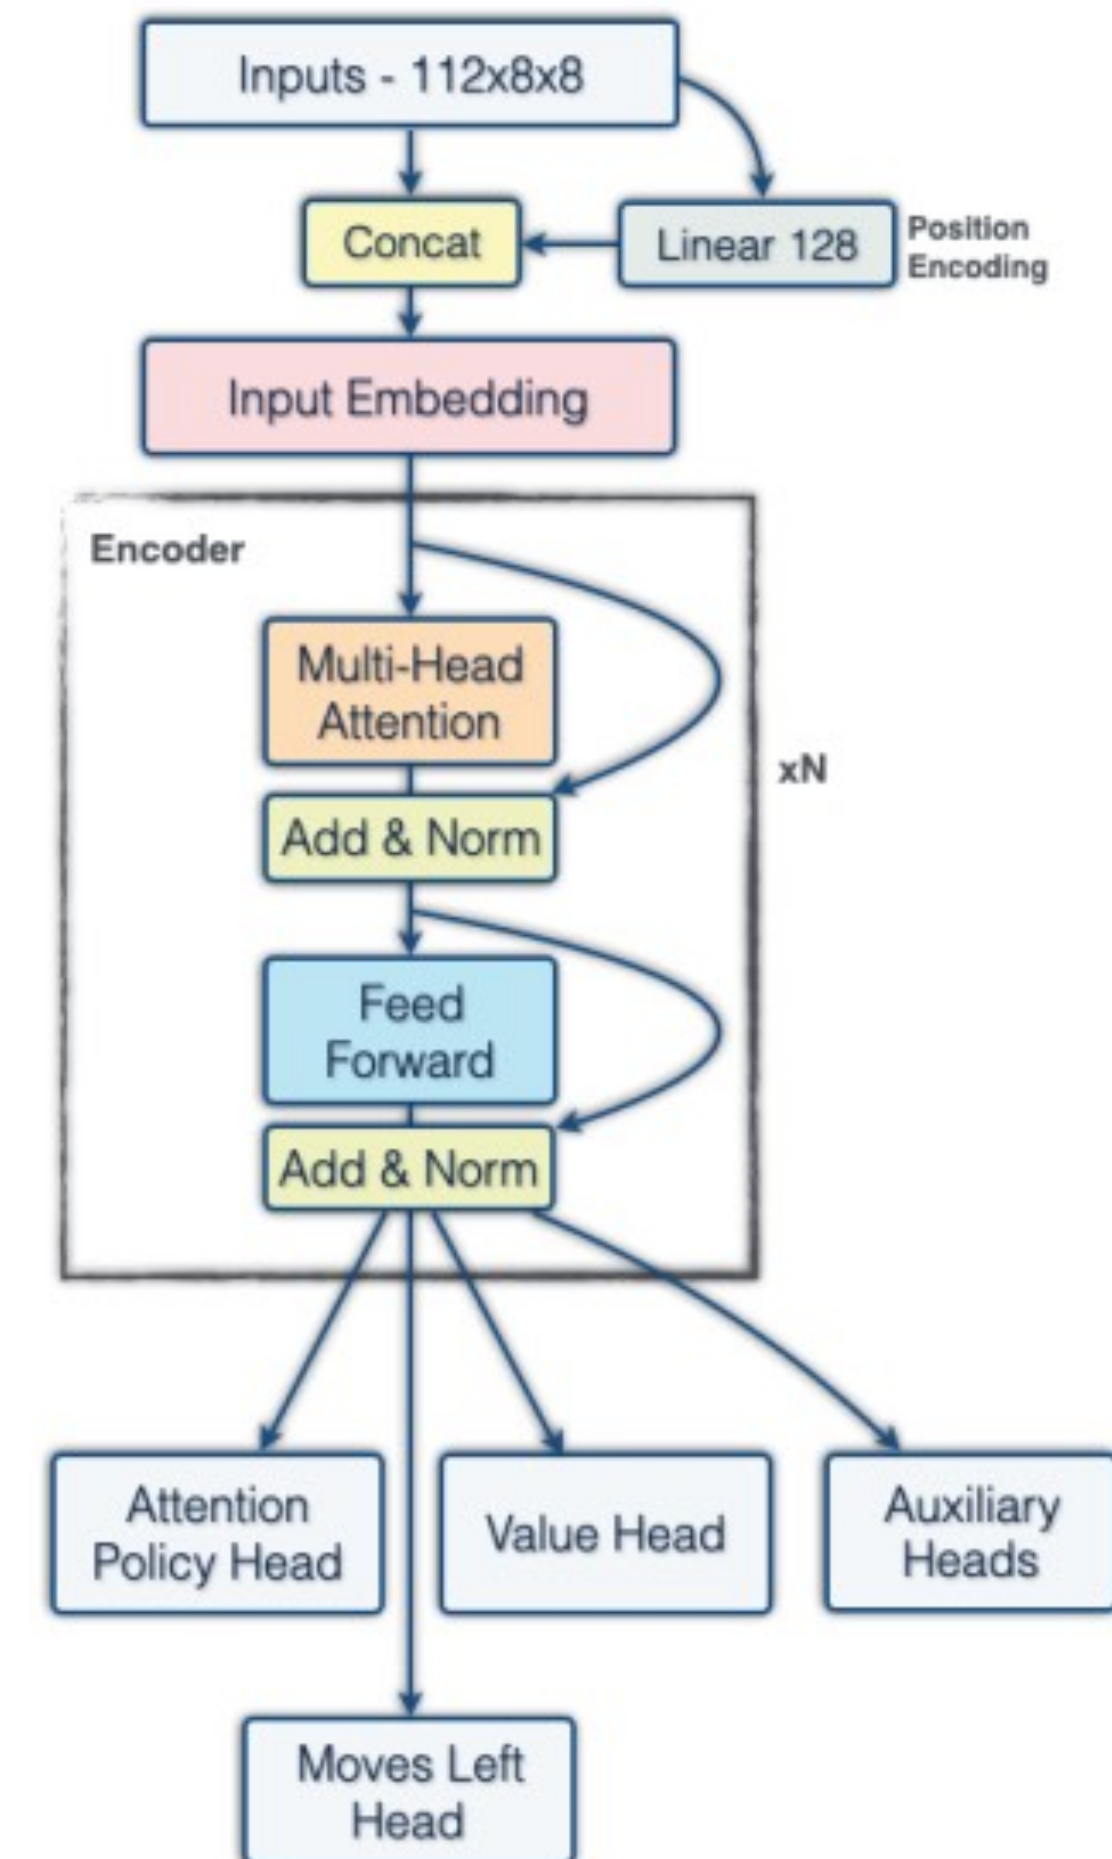
TRANSFORMERS AND BEYOND...?

LCZero convolutional models weak at finding long-range ideas on the chessboard

BT4 Transformer architecture —

64 tokens, 15 encoder layers, Embedding size 1024, FFN size 1536

— 191.3 million parameters (slightly smaller than convolutional models)



LEELA CHESS ZERO



TRANSFORMERS AND BEYOND...?

Chess-specific distance metrics
used in positional encoding
(instead of euclidean distance)

Dynamic positional information
based on nature of position -
smolgen model

Policy net for BT4 model **270 Elo**
stronger than that of T78 -
strongest convolutional model

