Playing with Memories – Instucen Trust

Digitising Play:
Using Artificial intelligence to Model and Reconstruct Ancient Games

Cameron Browne

Maastricht University

20/12/2020
Mumbai, Turkey

Maastricht University

Outline

1. Digital Ludeme Project
2. Modelling Games
3. Reconstructing Games
4. AI for Games

1. Digital Ludeme Project

Maastricht University

Evidence of Games

Lots of evidence of ancient board games:

- Boards, pieces, dice, etc.
- Last 5,000 years
- Most cultures worldwide

But almost never the rules!

Q. Can we use modern computational techniques to help improve our understanding of ancient games?

Maastricht University

Digital Ludeme Project

Five-year research project:
- Funded by the ERC (€2m)
- Maastricht University

1. **Model**
   Full range of traditional strategy games in a single playable digital database

2. **Reconstruct**
   Missing knowledge about ancient games

3. **Map**
   Spread of games throughout history

Maastricht University

European Research Council

Scope

Traditional strategy games

Traditional

- No proprietary owner
- Some historical longevity
- Connection with local culture

XII Scripta board from Laodicaea, Turkey

Strategy

- Reward mental skill
- Good decisions beat bad decisions

Model the 1,000 most “important” traditional strategy games

- Are documented
- Impact on evolutionary record of games

Timeline

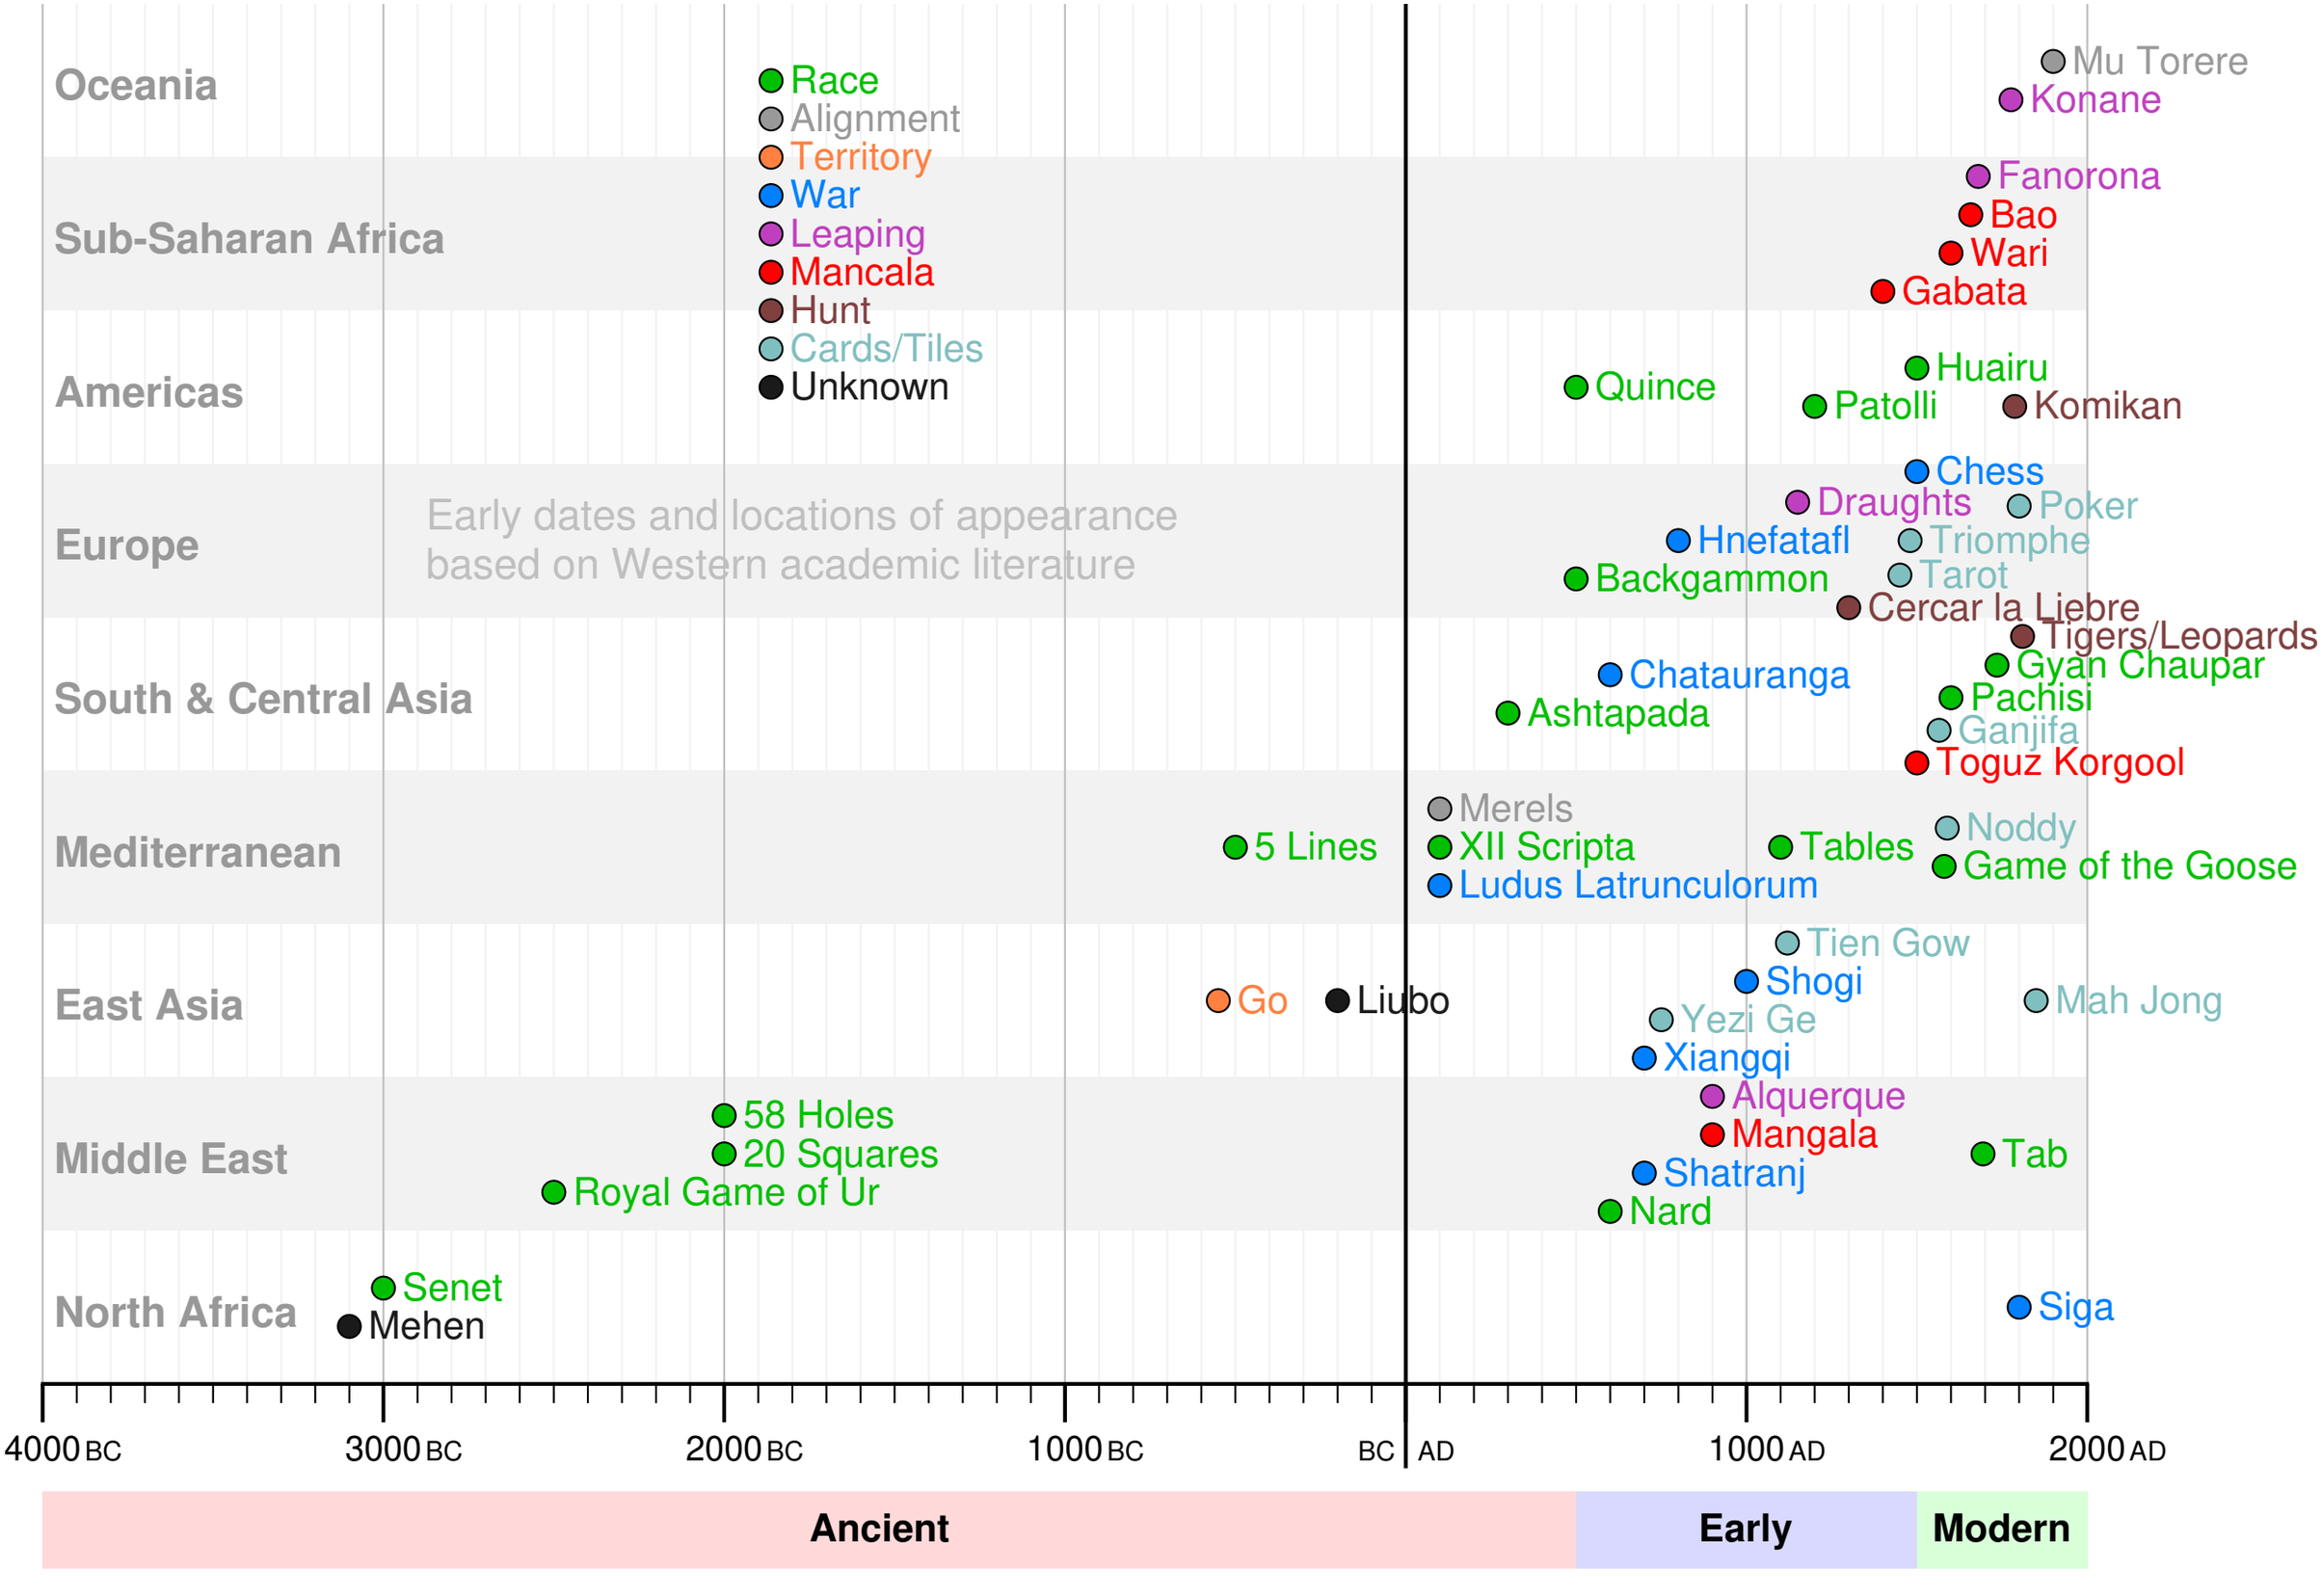
Oceania
Sub-Saharan Africa
Americas
Europe
South & Central Asia
Mediterranean
East Asia
Middle East
North Africa
Race
Alignment
Territory
War
Leaping
Mancala
Hunt
Cards/Tiles
Unknown
Early dates and locations of appearance
based on Western academic literature
Mu Torere
Konane
Fanorona
Bao
Wari
Gabata
Huairu
Quince
Patolli
Komikan
Chess
Draughts
Poker
Hnefatafl
Triomphe
Backgammon
Tarot
Cercar la Liebre
Tigers/Leopards
Gyan Chaupar
Chatauranga
Pachisi
Ashtapada
Ganjifa
Toguz Korgool
Merels
5 Lines
XII Scripta
Tables
Noddy
Game of the Goose
Ludus Latrunculorum
Tien Gow
Shogi
Go
Liubo
Mah Jong
Yezi Ge
Xiangqi
Alquerque
58 Holes
Mangala
20 Squares
Tab
Royal Game of Ur
Shatranj
Nard
Senet
Siga
Mehen
4000 BC
3000 BC
2000 BC
1000 BC
BC
AD
1000 AD
2000 AD
Ancient
Early
Modern

Timeline

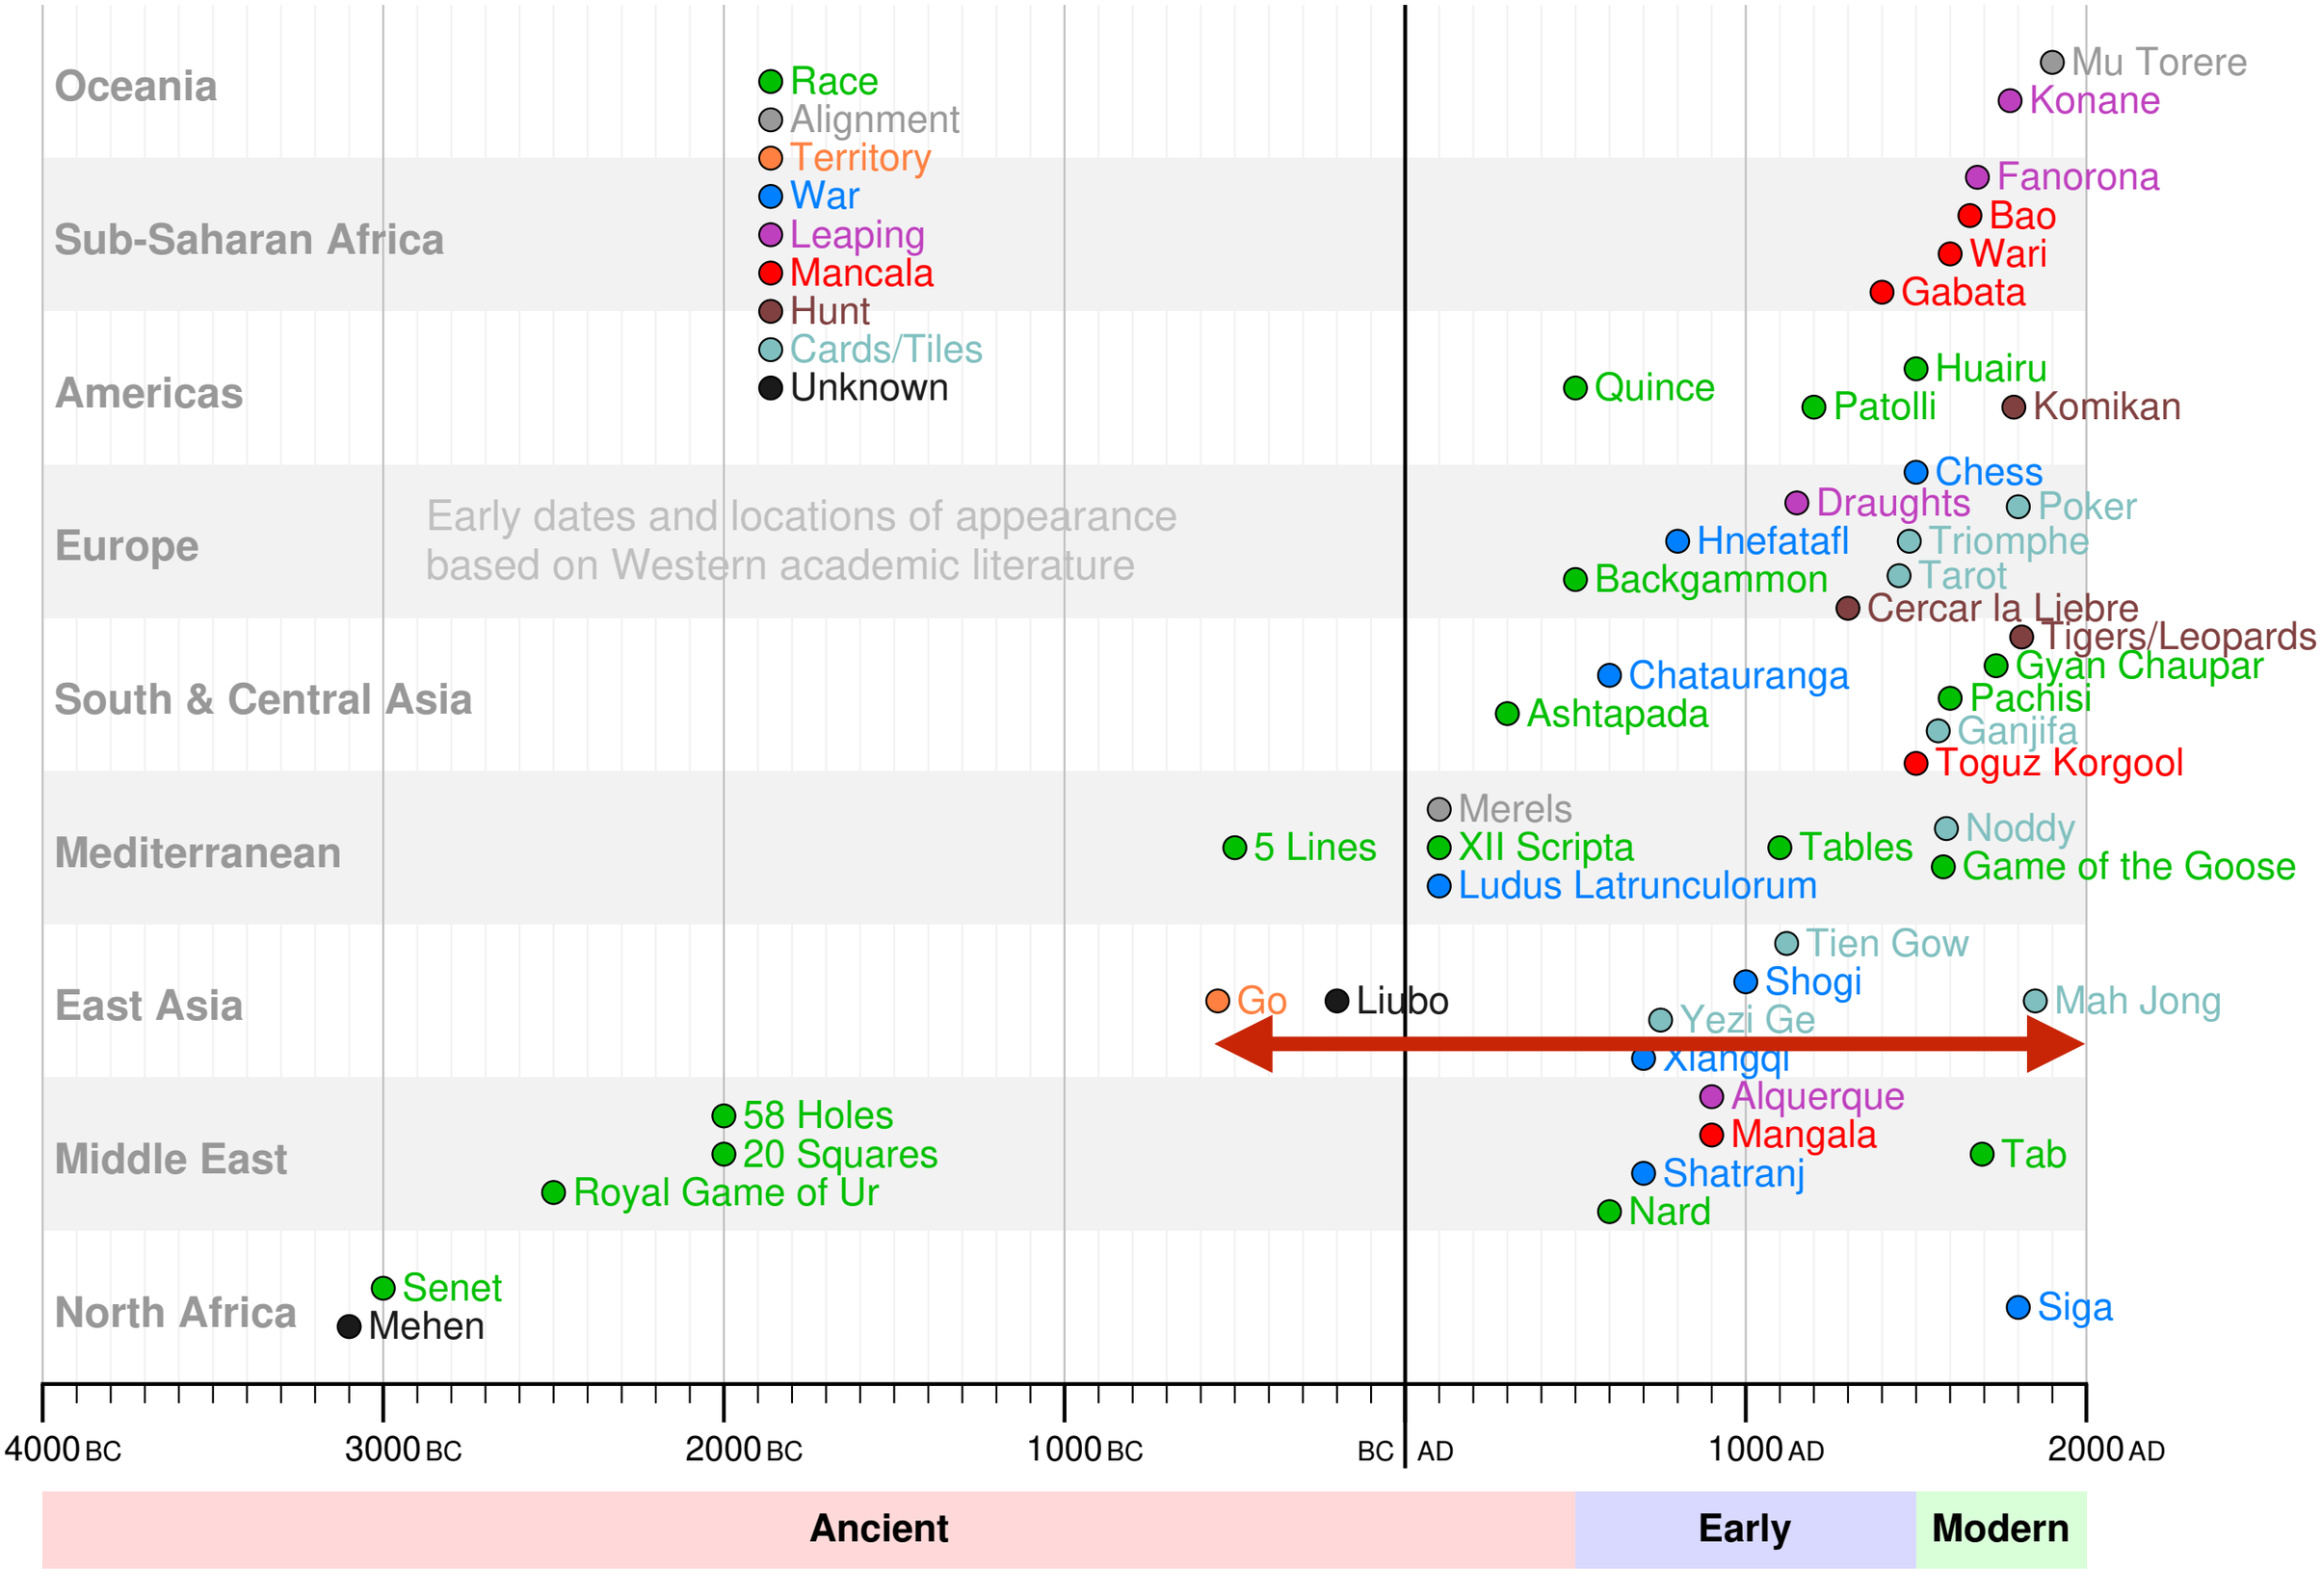
Oceania
Sub-Saharan Africa
Americas
Europe
South & Central Asia
Mediterranean
East Asia
Middle East
North Africa
Race
Alignment
Territory
War
Leaping
Mancala
Hunt
Cards/Tiles
Unknown
Early dates and locations of appearance
based on Western academic literature
Mu Torere
Konane
Fanorona
Bao
Wari
Gabata
Huairu
Quince
Patolli
Komikan
Chess
Draughts
Poker
Hnefatafl
Triomphe
Backgammon
Tarot
Cercar la Liebre
Tigers/Leopards
Gyan Chaupar
Chatauranga
Pachisi
Ashtapada
Ganjifa
Toguz Korgool
Merels
5 Lines
XII Scripta
Tables
Noddy
Game of the Goose
Ludus Latrunculorum
Tien Gow
Shogi
Go
Liubo
Mah Jong
Yezi Ge
Xiangqi
Alquerque
58 Holes
Mangala
20 Squares
Tab
Royal Game of Ur
Shatranj
Nard
Senet
Siga
Mehen
4000 BC
3000 BC
2000 BC
1000 BC
BC
AD
1000 AD
2000 AD
Ancient
Early
Modern

Timeline

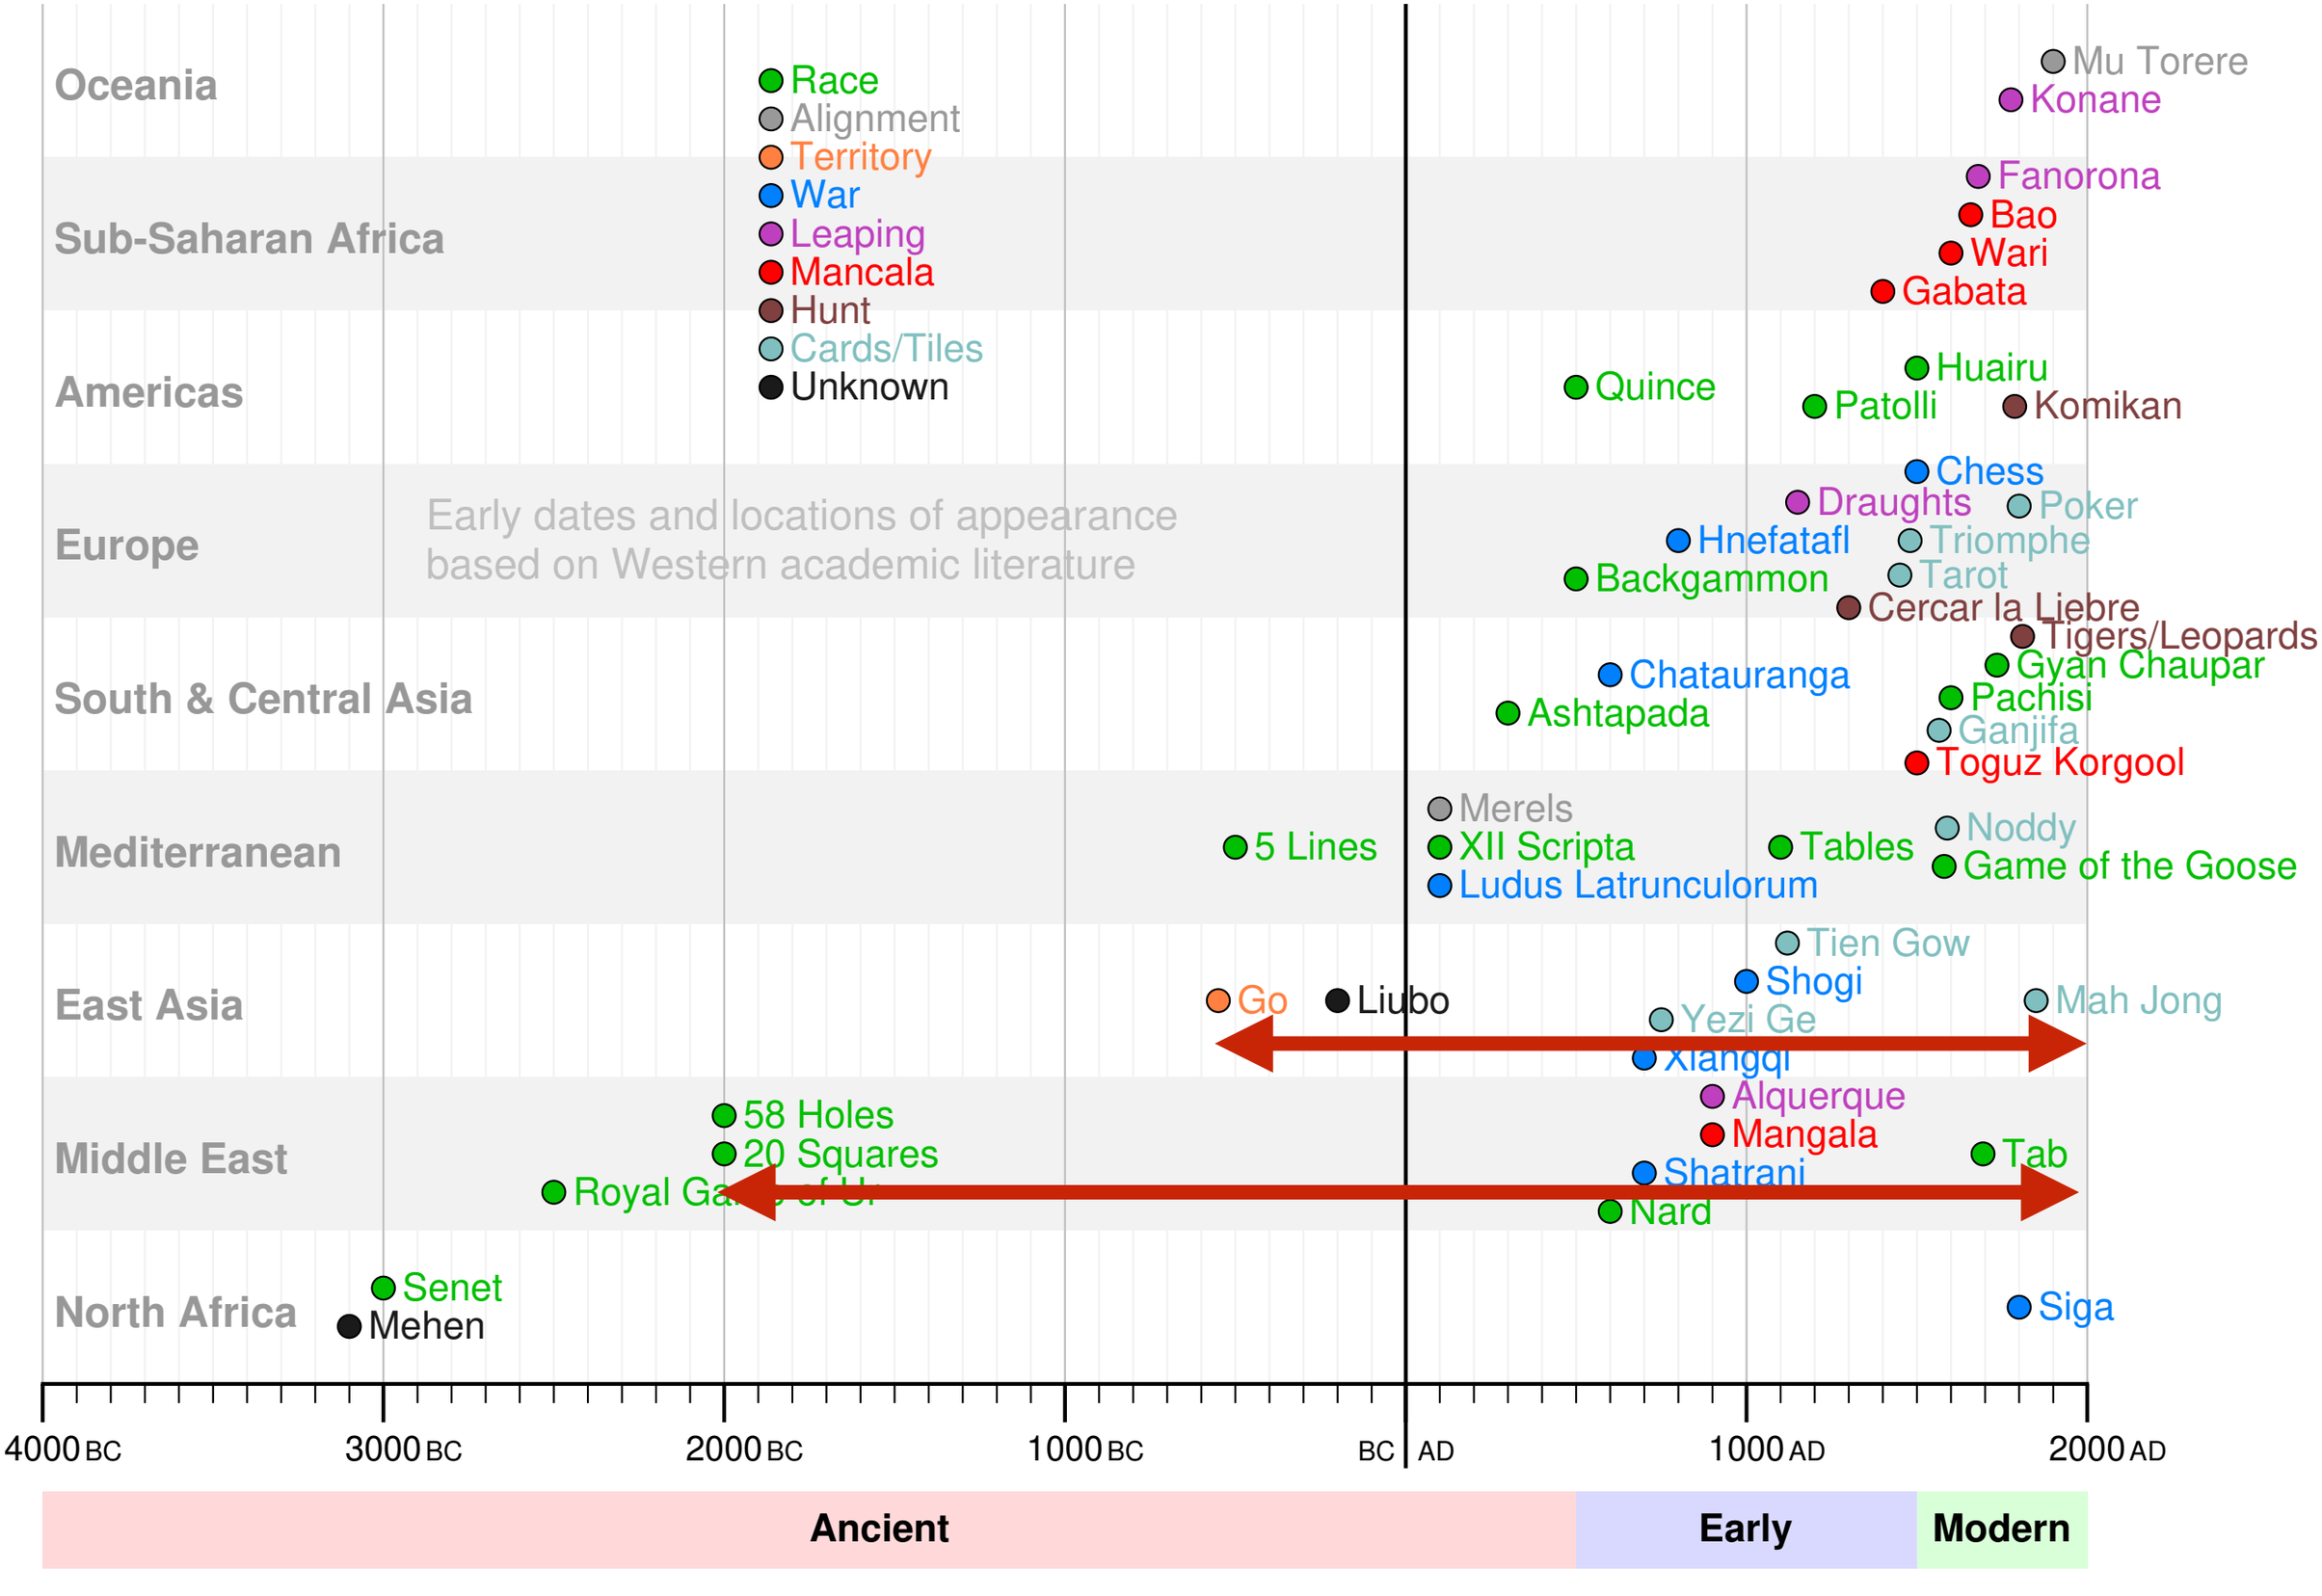
Oceania
Sub-Saharan Africa
Americas
Europe
South & Central Asia
Mediterranean
East Asia
Middle East
North Africa
Race
Alignment
Territory
War
Leaping
Mancala
Hunt
Cards/Tiles
Unknown
Early dates and locations of appearance
based on Western academic literature
Mu Torere
Konane
Fanorona
Bao
Wari
Gabata
Huairu
Quince
Patolli
Komikan
Chess
Draughts
Poker
Hnefatafl
Triomphe
Backgammon
Tarot
Cercar la Liebre
Tigers/Leopards
Gyan Chaupar
Chatauranga
Pachisi
Ashtapada
Ganjifa
Toguz Korgool
Merels
5 Lines
XII Scripta
Tables
Noddy
Game of the Goose
Ludus Latrunculorum
Tien Gow
Shogi
Go
Liubo
Mah Jong
Yezi Ge
Xiangqi
Alquerque
58 Holes
Mangala
20 Squares
Tab
Shatrani
Royal Ga
Nard
Senet
Siga
Mehen
4000 BC
3000 BC
2000 BC
1000 BC
BC
AD
1000 AD
2000 AD
Ancient
Early
Modern

2. Modelling Games

Maastricht University

Ludemes

A *ludeme* (“game meme”) is any game-related concept:

- Rules
- Equipment

e.g. `(tiling square)`

`(size 3)`

Atomic

Ludemes

A *ludeme* (“game meme”) is any game-related concept:

- Rules
- Equipment

e.g.

```
(tiling square)
```

```
(size 3)
```

```
(board
   (tiling square)
   (size 3)
)
```

Compound

Ludemes

A *ludeme* (“game meme”) is any game-related concept:

- Rules
- Equipment

e.g.

```
(tiling square)
```

```
(size 3)
```

```
(board
  (tiling square)
  (size 3)
)
```

```
(game "?"
  (players White Black)
  (board
    (tiling square)
    (size 3)
  )
  (move (add Own Empty))
  (end (win All (in-a-row 3)))
)
```

Ludemes

A *ludeme* (“game meme”) is any game-related concept:

- Rules
- Equipment

e.g.

```
(tiling square)
```

```
(size 3)
```

```
(board
  (tiling square)
  (size 3)
)
```

```
(game “Tic-Tac-Toe”
  (players White Black)
  (board
    (tiling square)
    (size 3)
  )
  (move (add Own Empty))
  (end (win All (in-a-row 3)))
)
```

Simple but powerful!

Data Sets

Three core data sets:

1. Games
   - Known rules, ludeme descriptions, classification, etc.
   - 600+ games (out of 1,000)

2. Ludemes
   - 500+ ludeme classes (move, from, hop, etc.)
   - 500+ ludeme constants (Left, Enemy, Empty, etc.)

3. Evidence
   - Artefacts, literature, artworks, ethnographies, etc.
   - 1,700+ entries so far

To do: Correlate dispersal of ludemes with evidence

Ludii

Software for performing the analysis

General game system
- Playing
- Analysing
- Generating

Currently >500 games

Free download:
- http://ludii.games

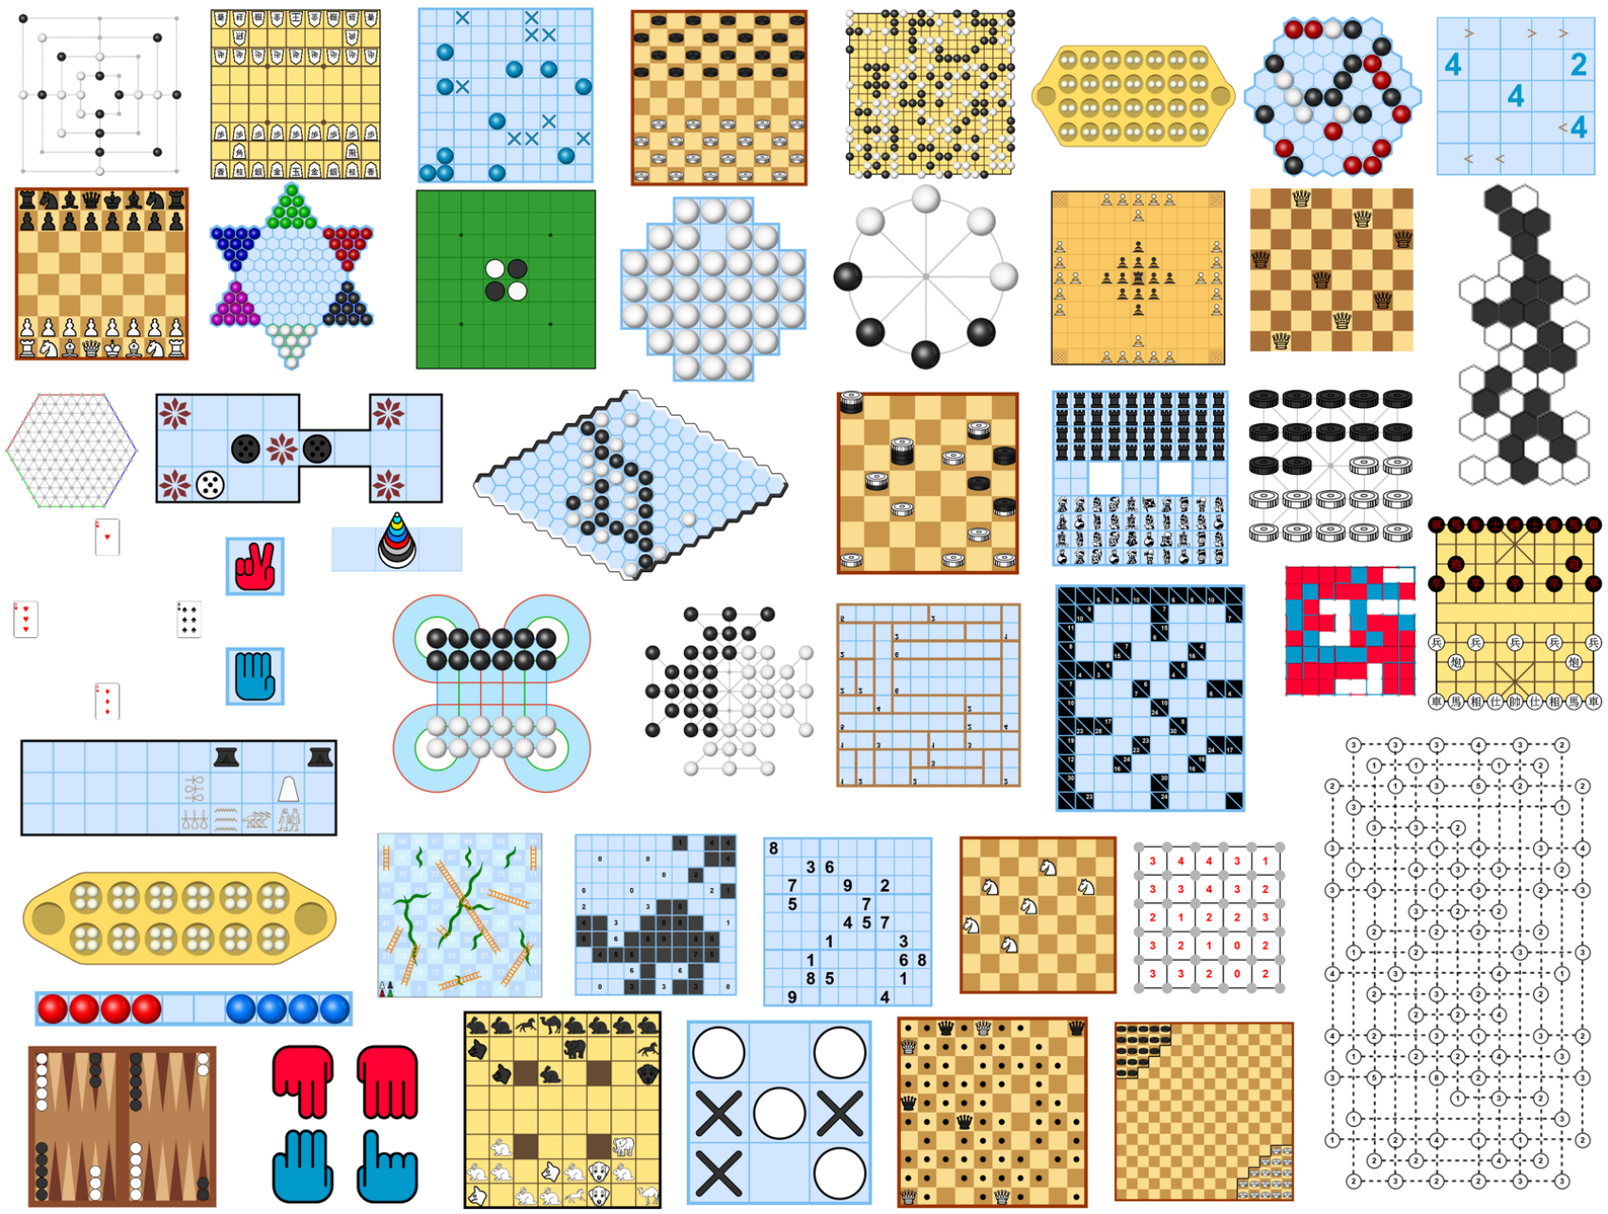

Maastricht University

3. Reconstructing Games

Reconstruction Task

Given:

- Partial knowledge of equipment and rules
- Historical/cultural context provided by evidence

We want to:

- Detect implausible reconstructions
- Suggest more plausible reconstructions

Example: Poprad Game

Poprad Game (Slovakia)
- Tomb dated to 375AD
- Germanic chieftain

Equipment
- 17x15/16 grid
- 2 x Colours
- 1 or 2 x Sizes?

Ulrich Schädler (2018)
- “An impossible task”
- Ludii may help!

Example: Poprad Game

Model in ludemic form:
- Describe known details
- Identify degrees of freedom

Not a single rule set!
- Distribution of rule sets

Improve the selection:
- Bias rules by confidence
- Filter out broken games

```
(game "Poprad"
    (players 2)
    (equipment {
        (board (rectangle 17 [15 | 16]))
        [
            (piece "Disc" P1)
            (piece "Disc" P2)
            |
            (piece "Disc" P1)
            (piece "Disc1" P2 value:1)
            (piece "Disc2" P2 value:2)
            |
            (piece "Disc1" P1 value:1)
            (piece "Disc2" P1 value:2)
            (piece "Disc1" P2 value:1)
            (piece "Disc2" P2 value:2)
        ]
    })
    (rules
        [(start [*]) | *]
        (play [(move Add (to (sites Empty))) | *])
        (end [*])
    )
)
```

Evaluating Reconstructions

1. Historical Plausibility

- Is rule set compatible with historical/cultural context?
- Are concepts contemporary?

2. Game Quality

- Does it play well?
- Is it interesting?
- Is it likely to be transmitted?

Measuring Game Quality

How do we actually measure this?

Computational Creativity provides a framework (Ritchie 2007):

1. Novelty

- Rule set is different to existing rule sets

2. Typicality

- Rule set provides a playable game

3. Quality

- Rule set provides an interesting game

Typicality

Definition of a game (*based on Rules of Play*, 2003):

> ***A game is a fair contest with achievable outcomes.***

A rule set is ***typical*** if it is:

1. Well-Formed
- Conforms to the Ludii grammar

2. Executable
- Compiles to executable (Java) bytecode

3. Correct
- Runs without error

4. Playable
- Allows all players to make at least one move

5. Gamelike
- Provides a fair contest with achievable outcomes

Gamelike

A rule set is *gamelike* if it is:

1. Balanced

- All players win

2. Decisive

- Most games produce a result (win/loss)

3. Good Length

- Not too short or too long

Reliable, easy, fast to detect

- Quickly eliminate large numbers of flawed games

Example: Mu Torere

Mu Torere (New Zealand, 18thC):

- Living players
- Full knowledge of rules

Move a piece of your colour
to the adjacent empty point,
<u>if it is next to an enemy piece</u>

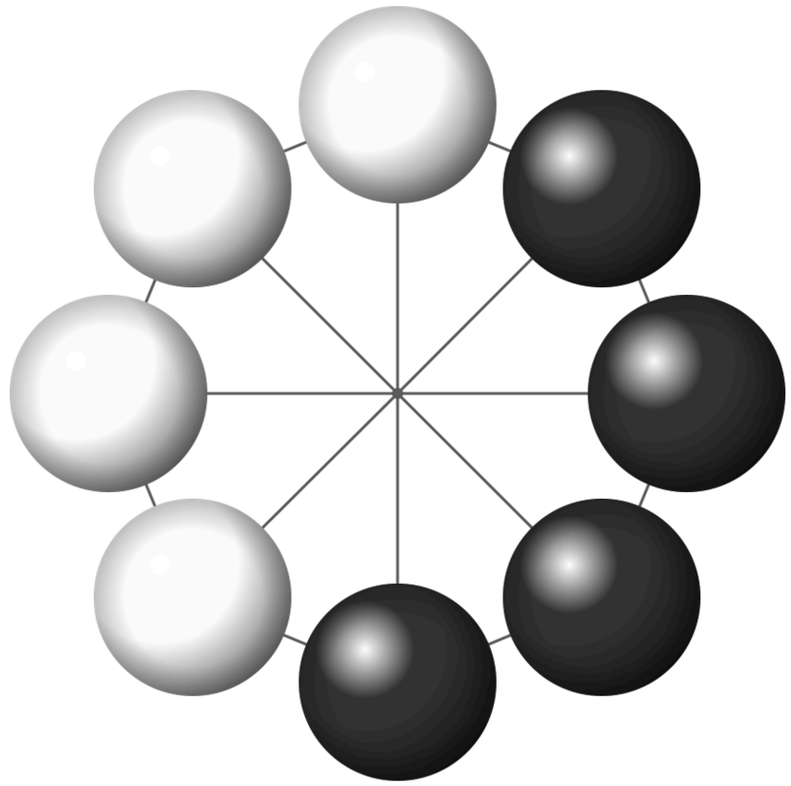

Example: Mu Torere

Mu Torere (New Zealand, 18thC):
- Living players
- Full knowledge of rules

Move a piece of your colour
to the adjacent empty point,
<u>if it is next to an enemy piece</u>

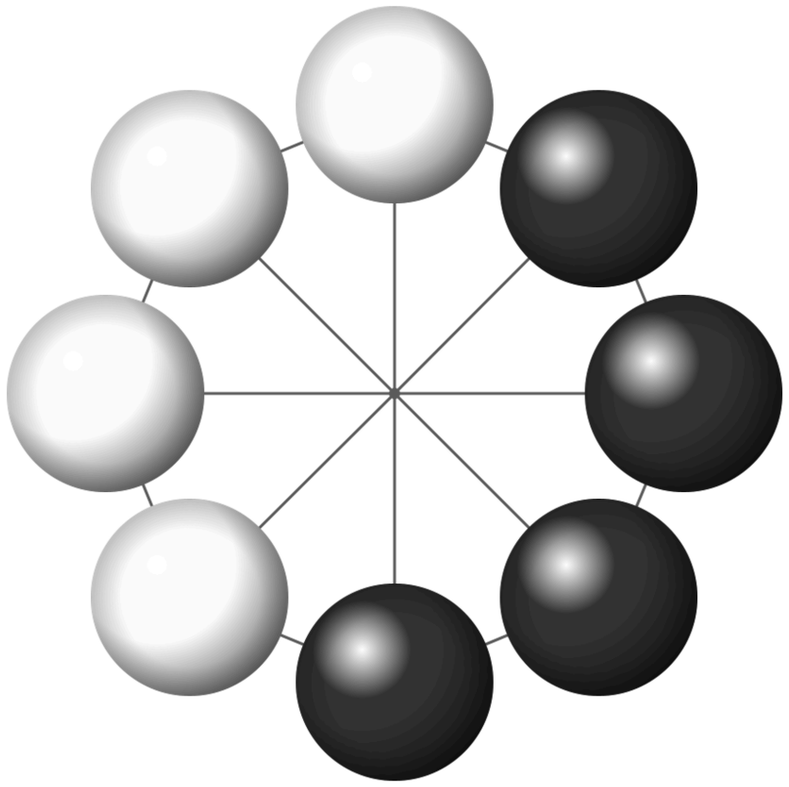

Marcia Ascher (1987) survey:
- Two historical accounts forget this rule
- Win on first move:
  - Unbalanced
  - Decisive
  - Bad game length

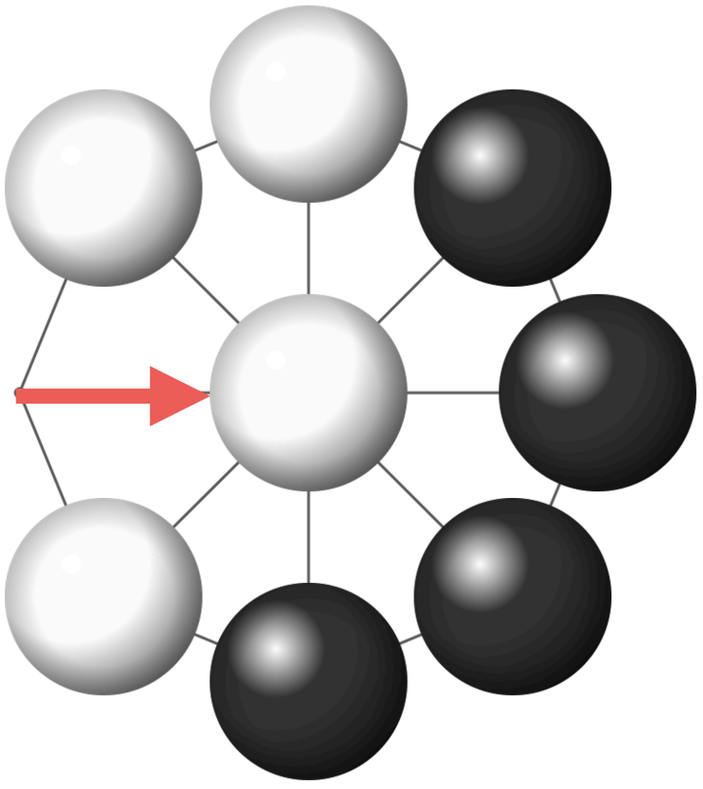

Game Quality

If pass **typicality** test then measure for **quality**

- Much harder!

Criteria:

- Strategic depth
- Uncertainty
- Drama
- Tension
- Clarity
- Skill/chance tradeoff

Game Quality

If pass **typicality** test then measure for **quality**

- Much harder!

Criteria:

- **Strategic depth**
- Uncertainty
- Drama
- Tension
- Clarity
- Skill/chance tradeoff

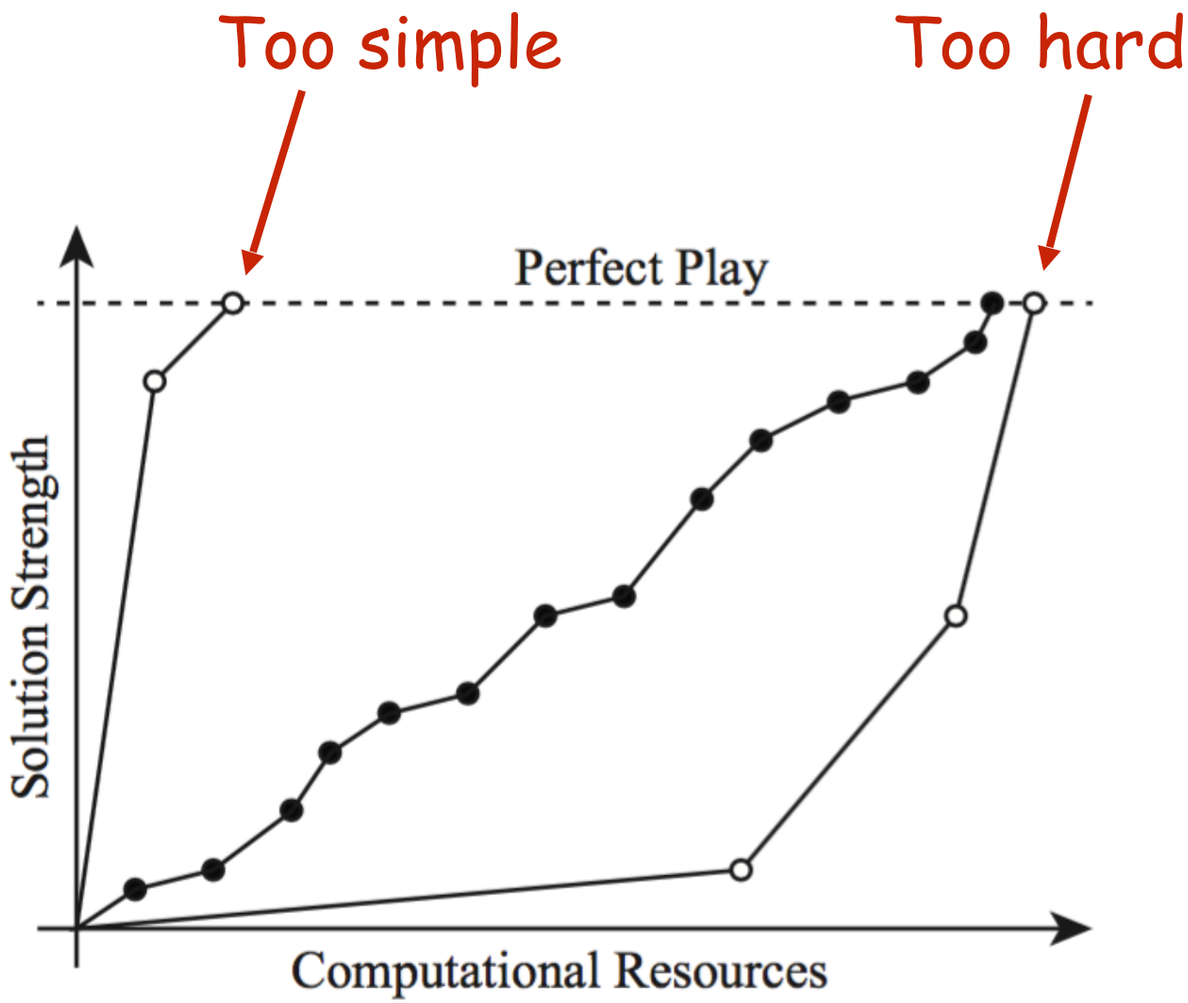

Strategy Ladder
Lantz *et al.* (*AAAI'17*)

4. AI for Games

Automated Playtesting

Get AI agents to play against each other (“AI self-play”)
1. AI agents A and B play 100 games
2. Check for typicality
3. Measure quality

Time vs reliability:

Weak (random) AI agents
- Milliseconds per move
- Unreliable results in seconds

Strong(er) AI agents
- Seconds per move
- Reliable results in hours or days

AI Approaches

1. Tree-Based Methods (1950s)

- Exhaustive search
- Requires heuristic knowledge

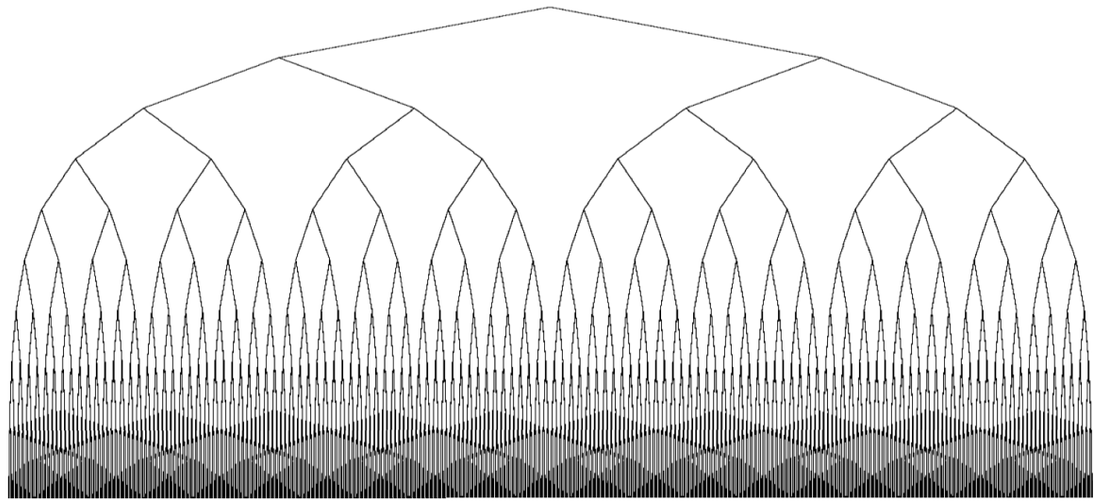

2. Monte Carlo Methods (1930s)

- Random sampling
- No heuristic knowledge

3. Monte Carlo Tree Search (2007)

- Build tree from random sampling
- Revolutionised game AI

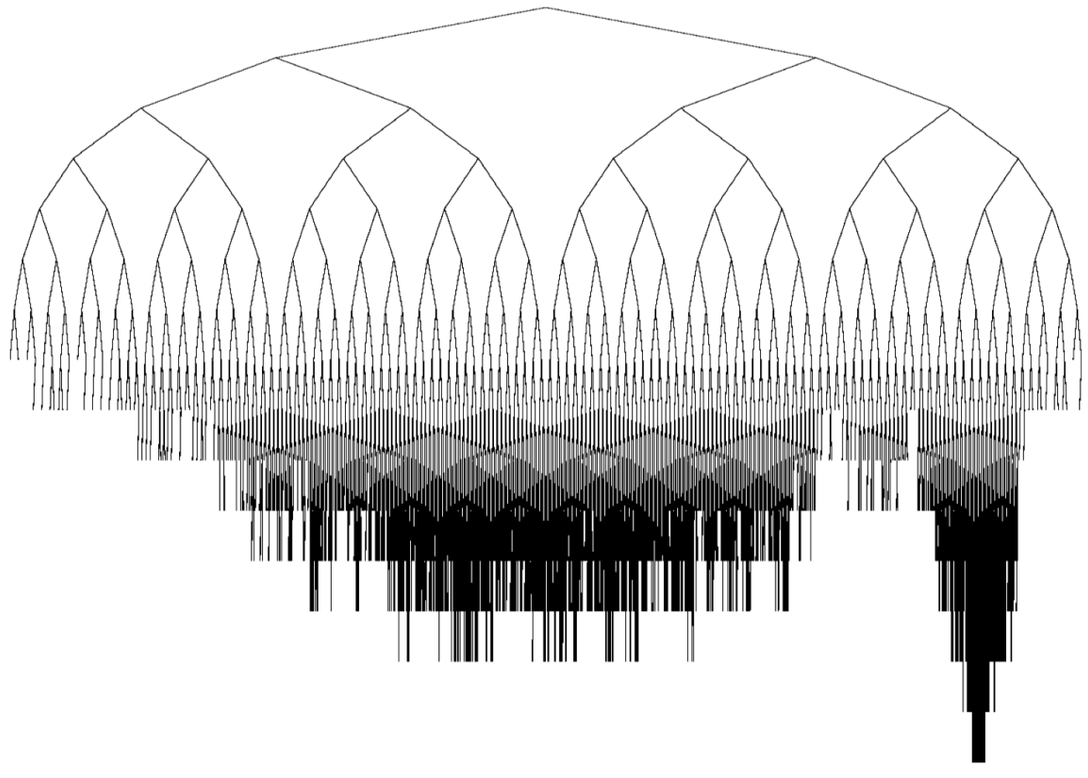

AI Approaches

4. Deep Learning (2016)

- MCTS with neural networks
- Superhuman playing strength
- Hugely expensive

Timeline:

- AlphaGo beats Lee Sedol 4–1 (2016)
- AlphaGo Zero beats AlphaGo 100–0 (2017)
- AlphaZero learns Go, Chess, Shogi (2017)

- Removing human expert knowledge improved strength!

AI Strength

Stephen Tavener (2020)

- 3x3 mini-game experiment

3x3 Go

- Win for P1
- More search = stronger result

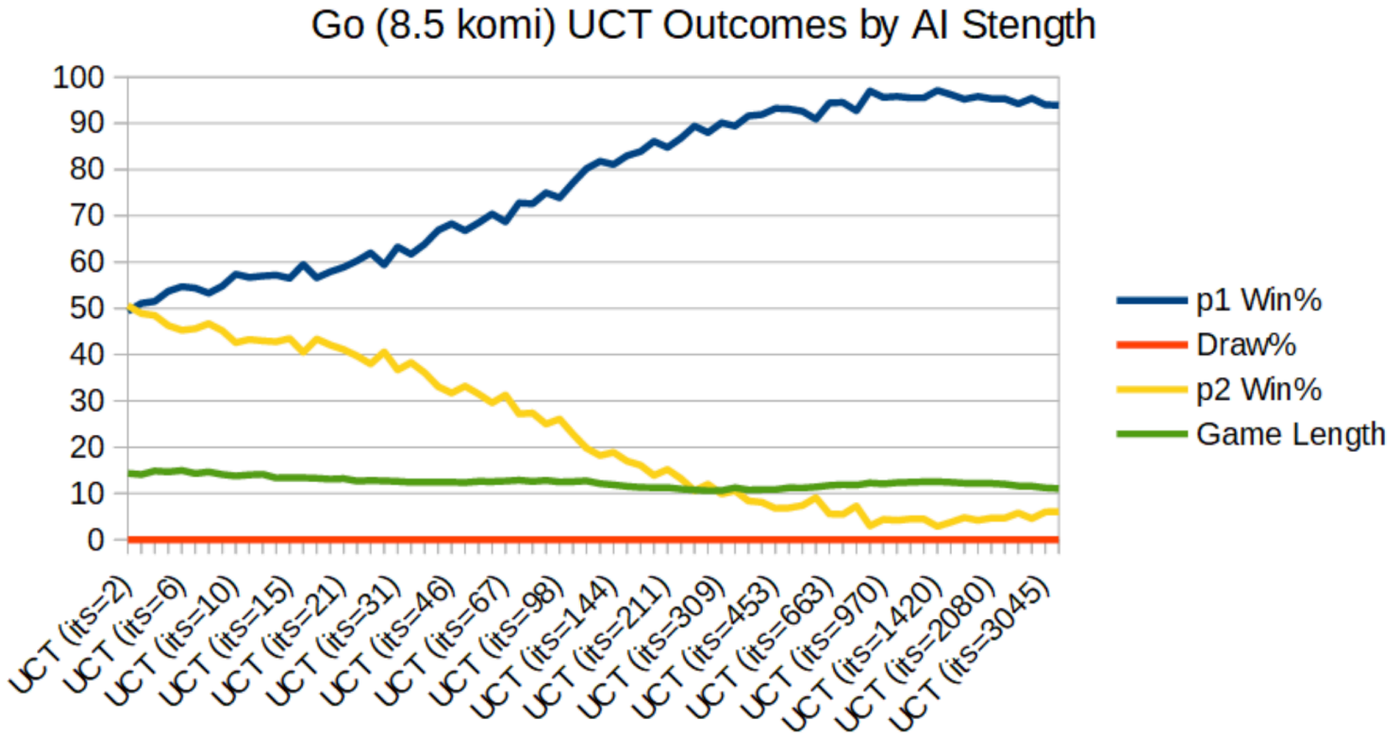

AI Strength

Stephen Tavener (2020)
- 3x3 mini-game experiment

3x3 Go
- Win for P1
- More search = stronger result

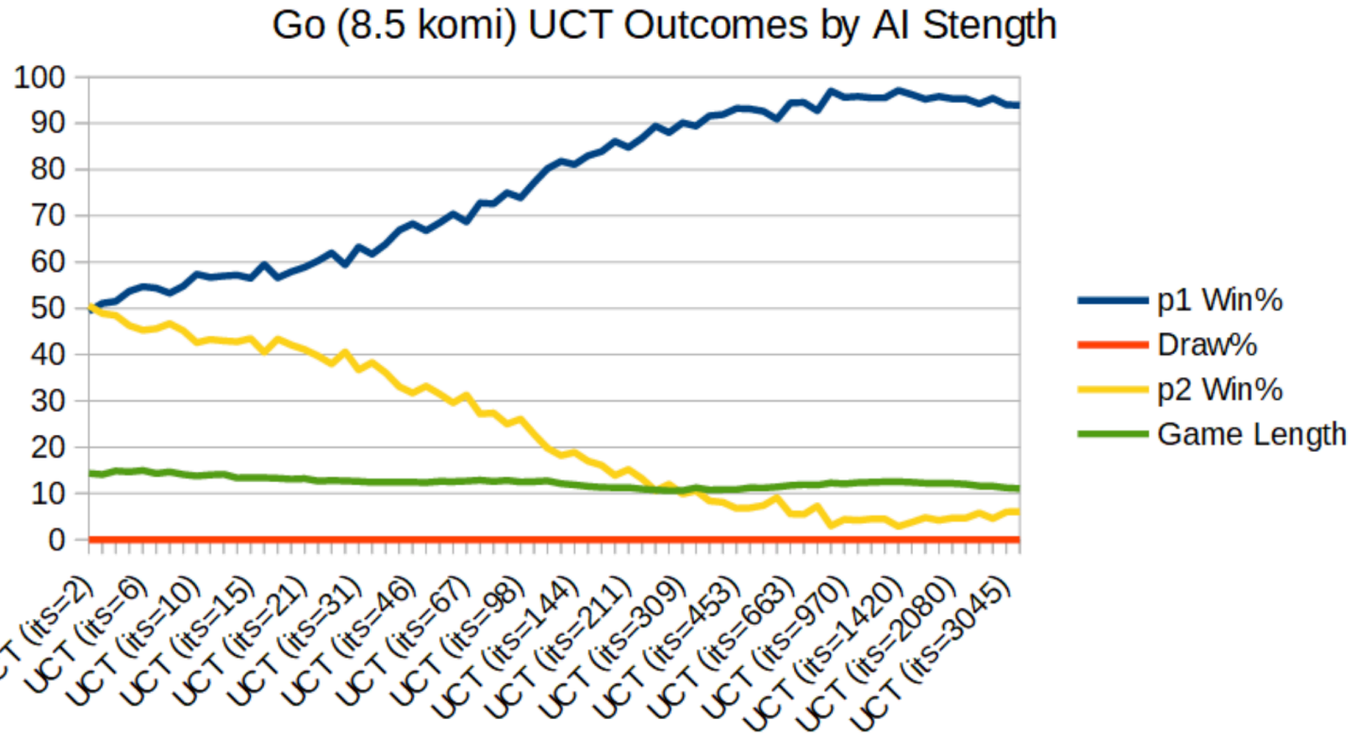

3x3 Chessline
- Behaviour changes based on AI strength
- Winning strategy for P1

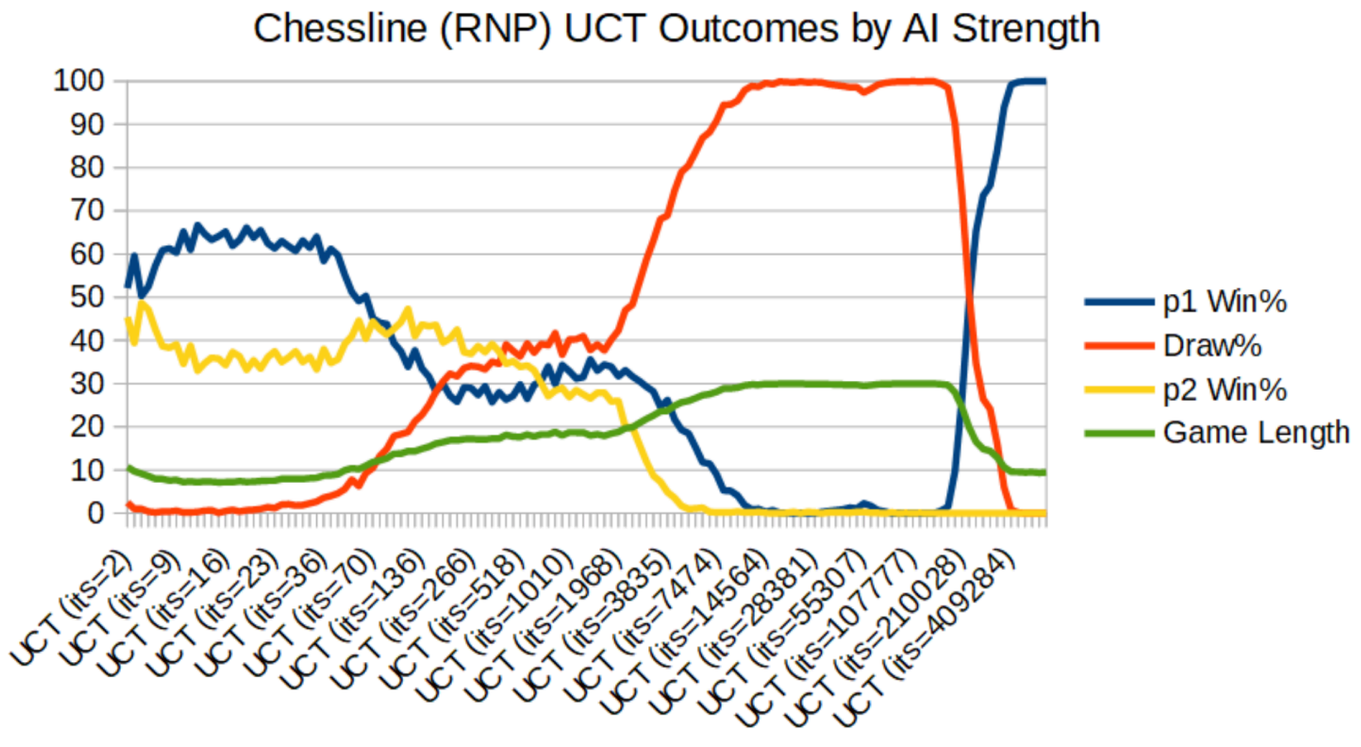

Example: Hnefatafl

Hnefatafl “Viking Chess”
- Scandinavia (c.800AD)
- No original rules found

Linnaeus (1732)
- Saw Tablut, transcribed rules (in Latin)

Smith (1811)
- Translated into English

Murray (1913) *History of Chess*
- Assumed same rules for Hnefatafl
- Became de facto

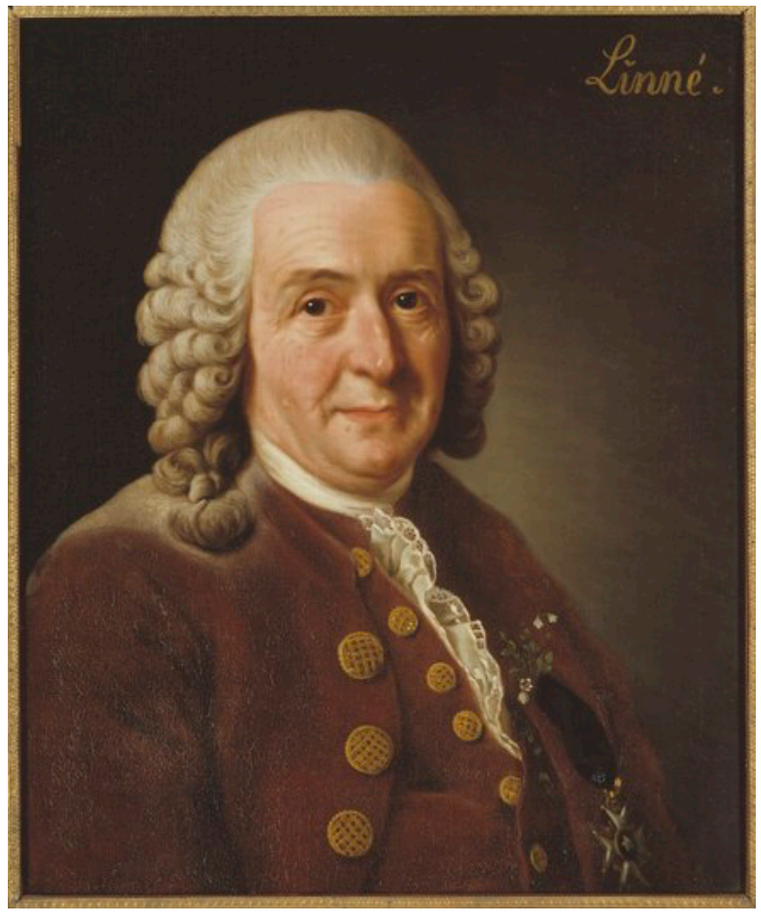

Carl Linnaeus (1707–1778)

Example: Hnefatafl

BUT…

Smith made a
bad translation of
the king capture rule

Original Latin

- “*likewise the king*”
- Flanked
- Easy

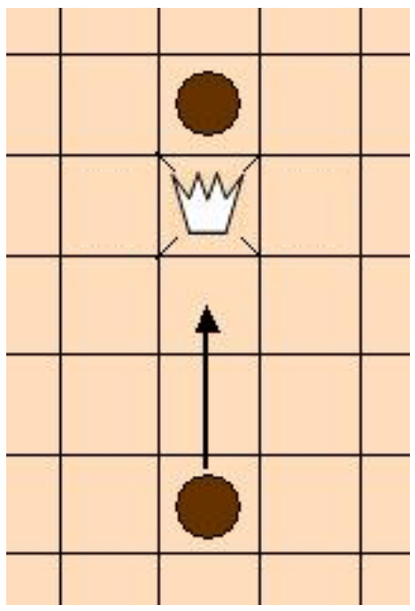

Smith’s Version

- “*except the king*”
- Surrounded
- Hard!

[DEMO]

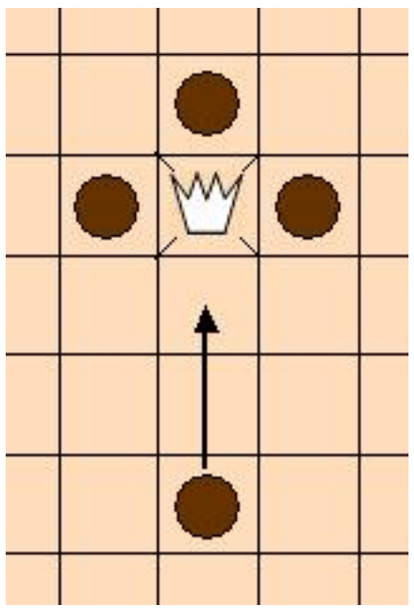

But There's More…

An intelligent player should find a winning strategy:
1. Form a ring
2. Constrict

Bias swings towards attackers

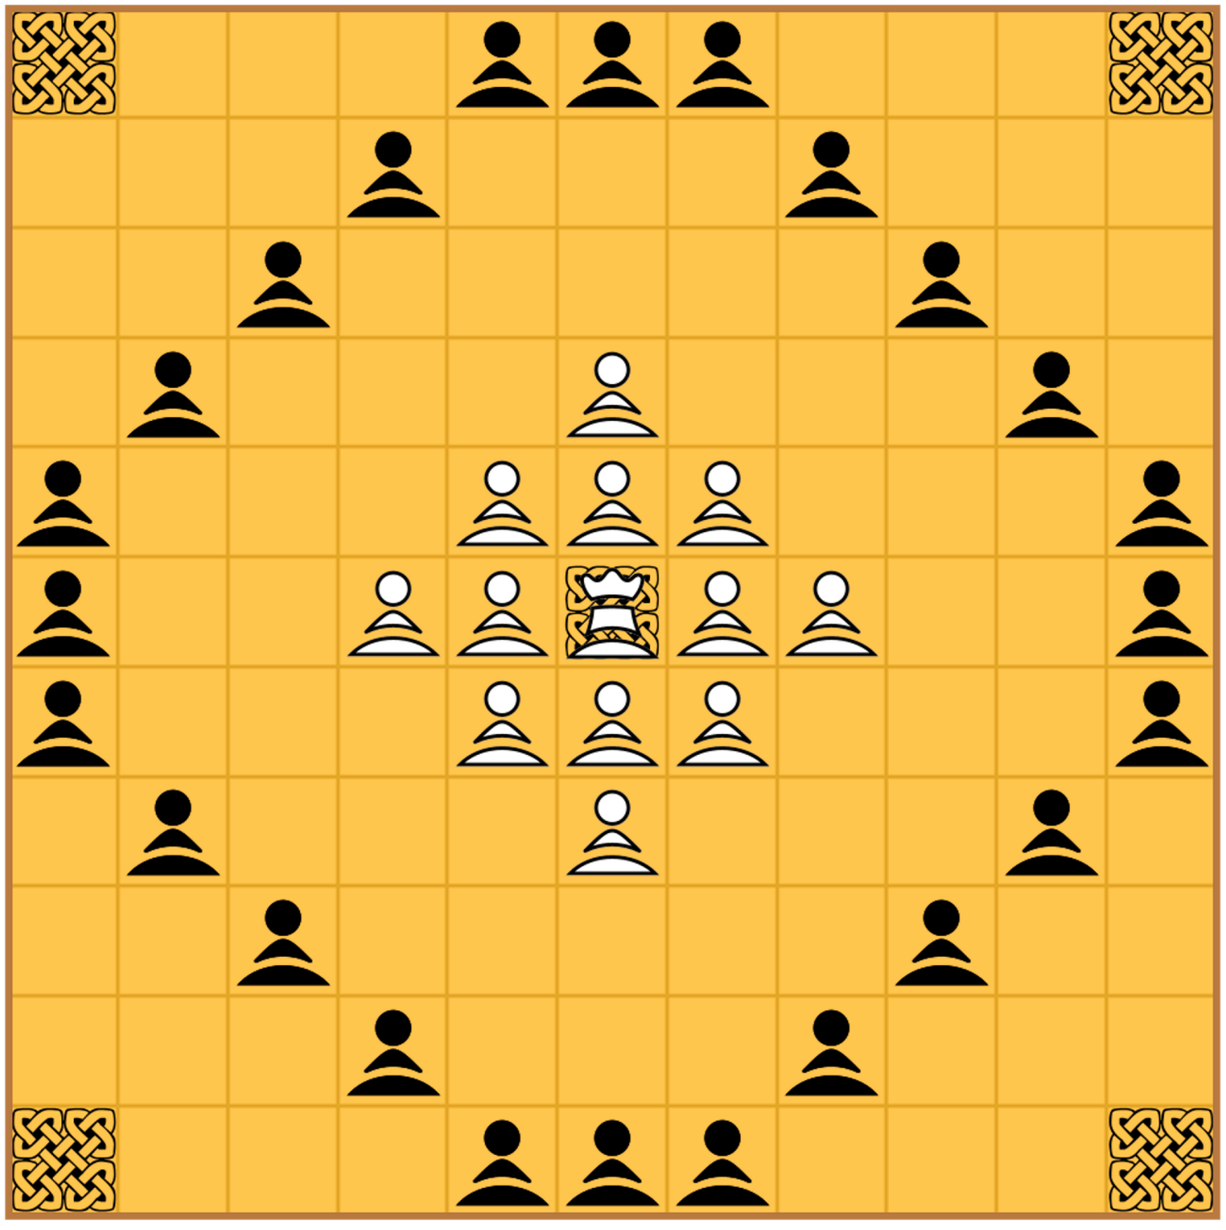

But There's Even More…

An even more intelligent player should find a spoiling strategy:

- Make "fortress"
- Move king back and forth

Neither player can win

Copenhagen Rules

Where to pitch AI level?

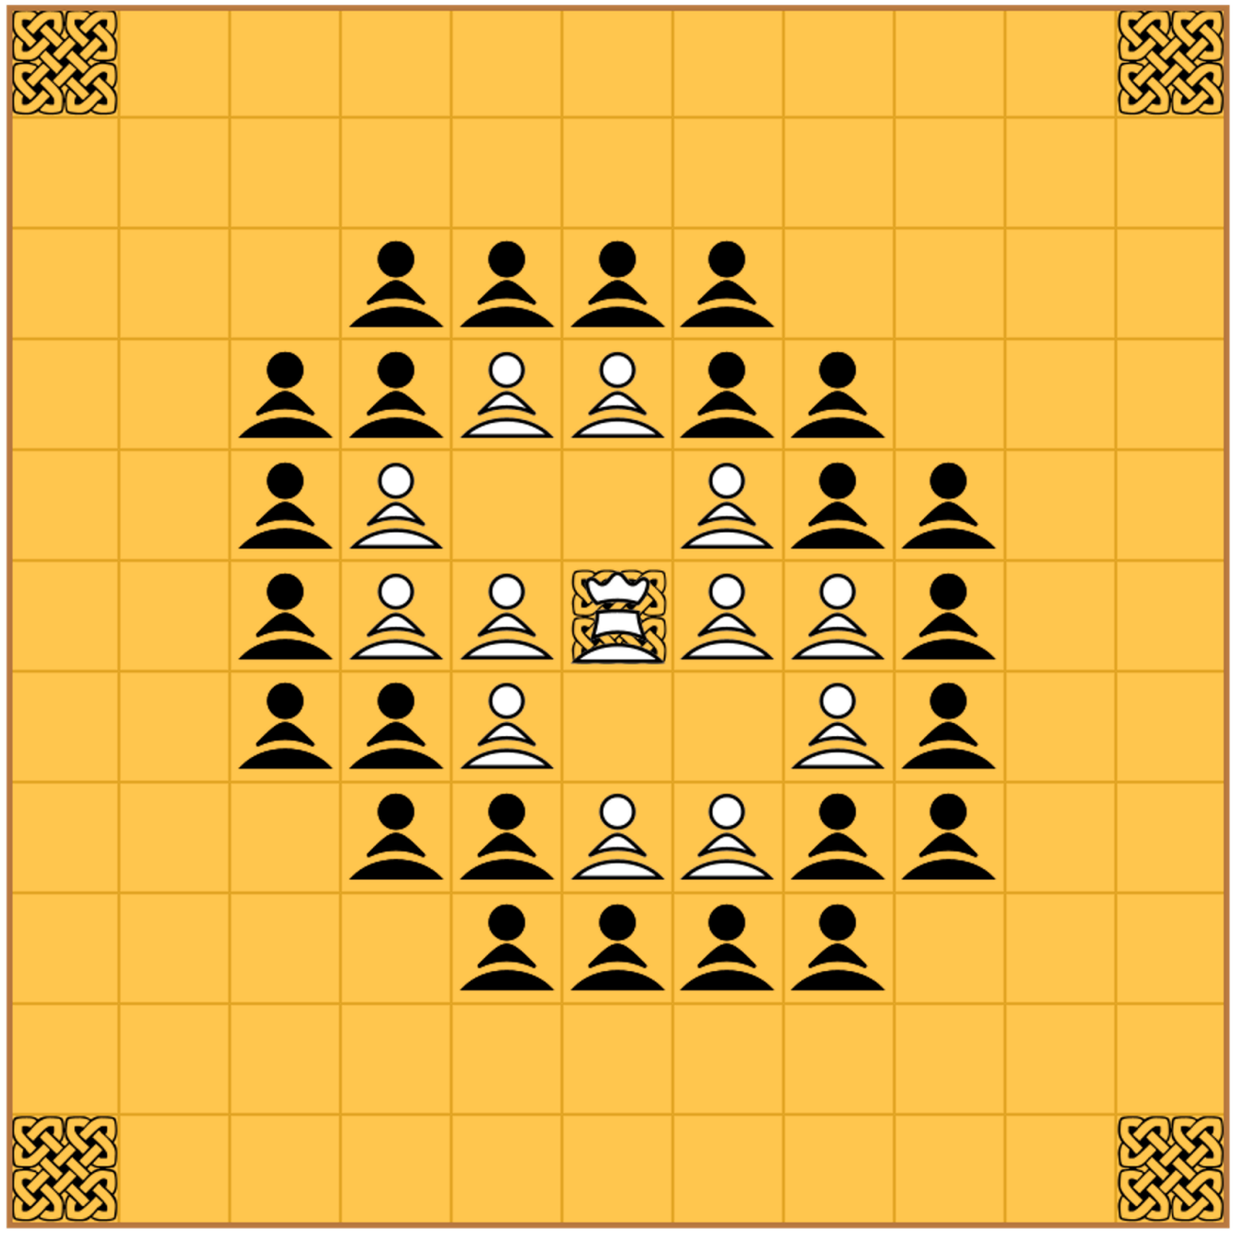

Human-Level AI

We don’t want superhuman AI!
- Draughts and Chess drawish at world champion level
- Not the average human experience

We don’t want random agents:
- Not the average human experience

“Human-level AI”:
- Win 50% of games against top 50% of players

Artificial Stupidity

We actually need to weaken the AI in some cases!

e.g. Taikyoku Shogi (Japan, 15thC)

- Most complex board game played by humans
- 402 pieces each (209 types)

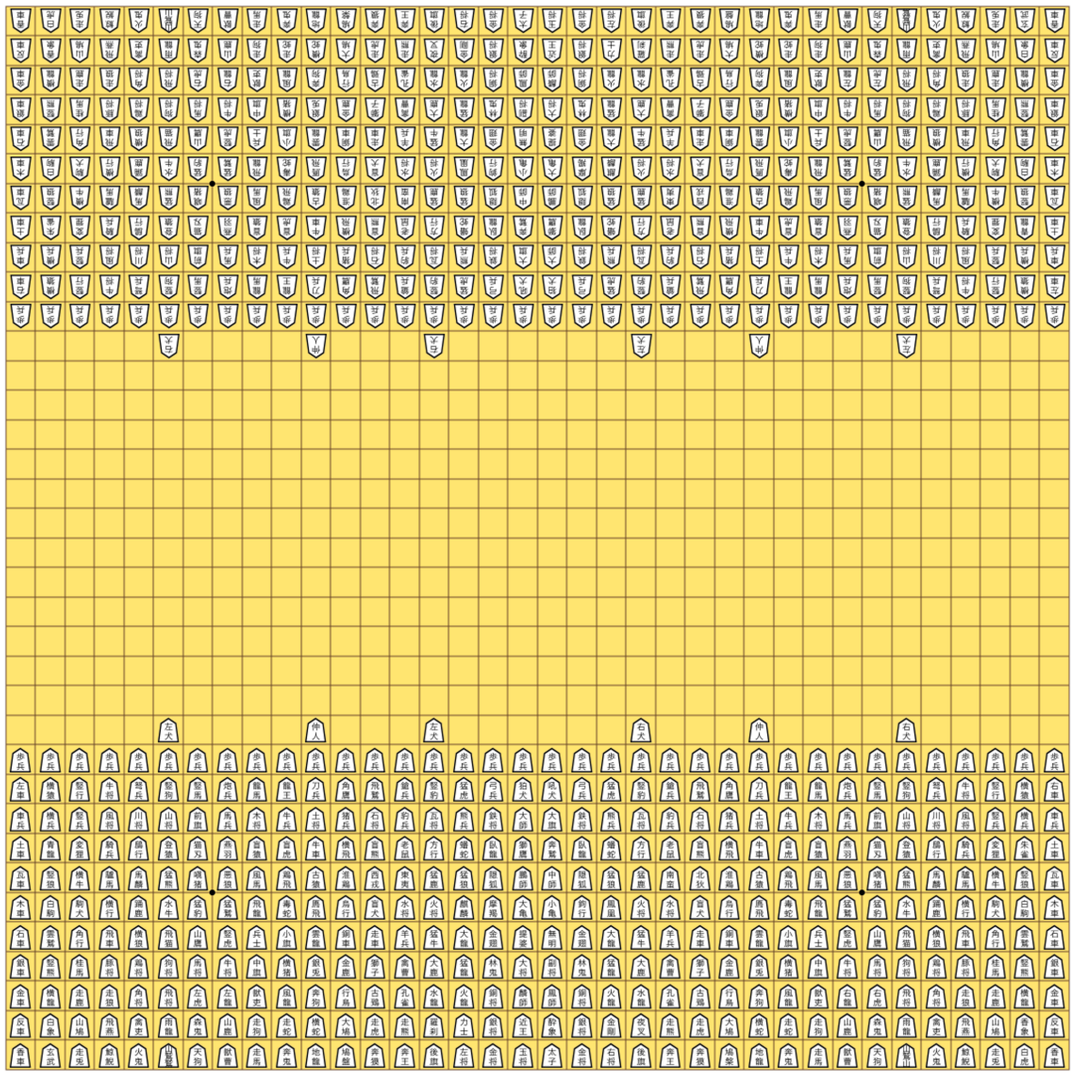

1-ply lookahead will beat any human

Must actually hobble the AI

Conclusion

Thank You

Questions?

Digital Ludeme Project

http://ludeme.eu

http://ludii.games

European Research Council

Established by the European Commission

Maastricht University