Learning and Adaptation of Game AI

Pieter Spronck

Universiteit Maastricht / MICC-IKAT

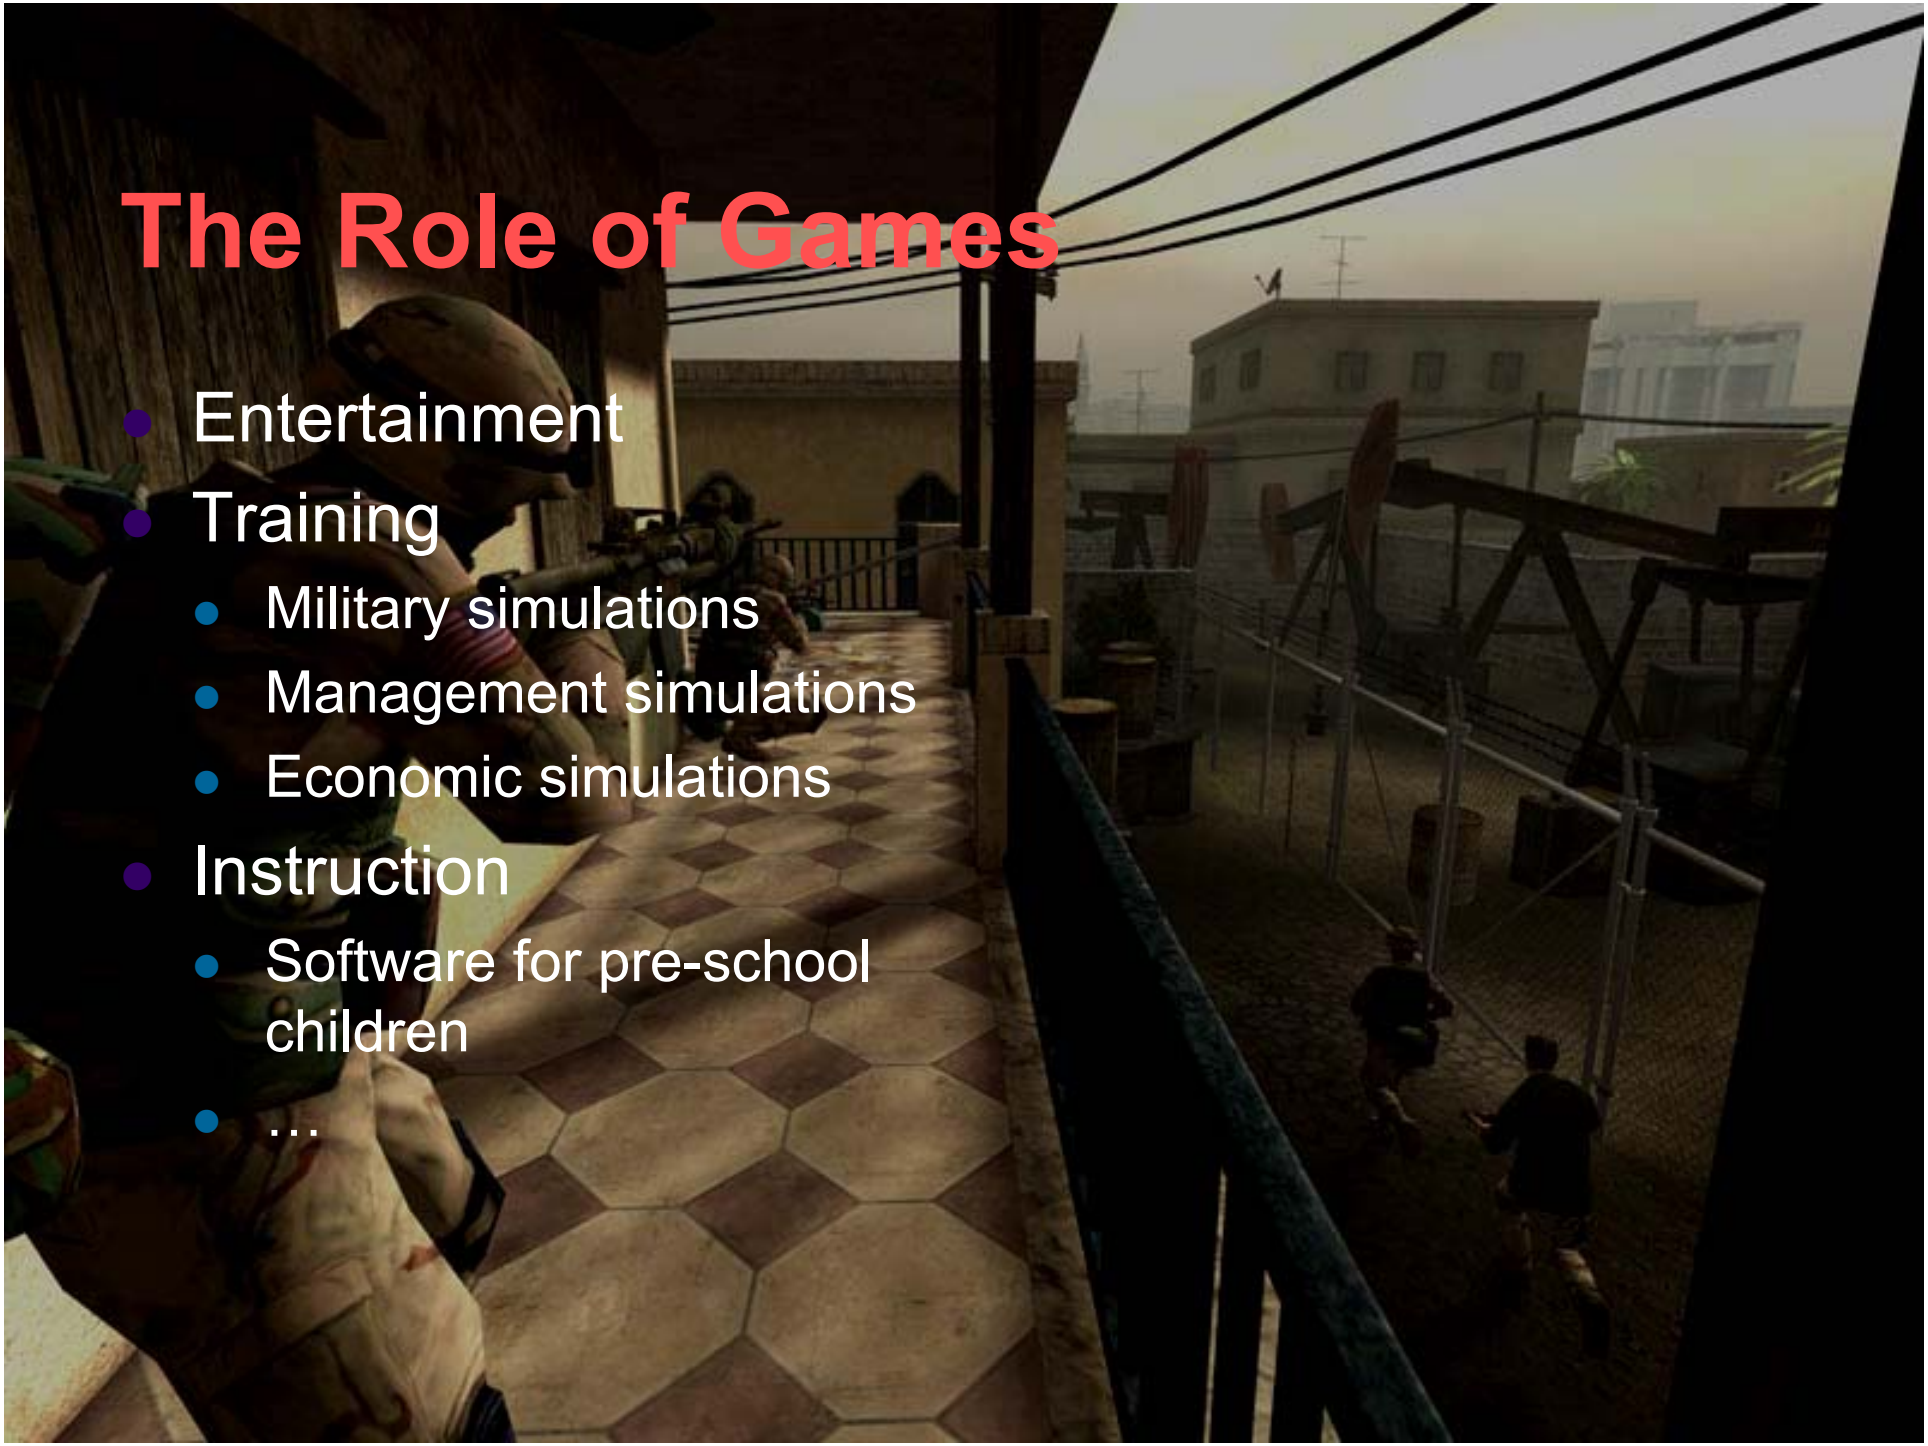

The Role of Games

- Entertainment
- Training
  - Military simulations
  - Management simulations
  - Economic simulations
- Instruction
  - Software for pre-school children
  - …

Learning and Enjoying Chess

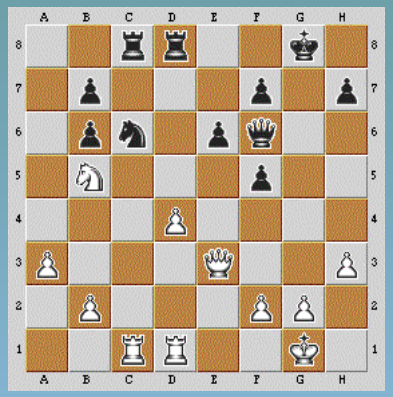

Deep Blue (1997)

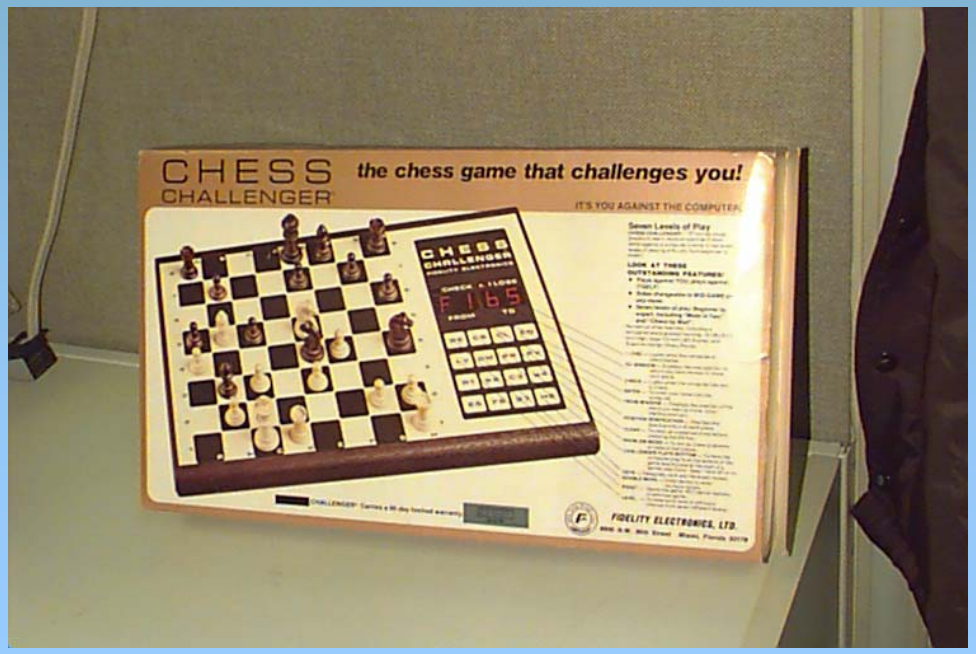

Chess Challenger (1978)

Learning and Enjoying Games

- Computer should *be able* to play strongly
- Computer should *adapt* to the level of skill of the human player
- Computer should constantly offer *new challenges*

In short: Computer and human increase their playing skill in *parallel*

- Manual
- Coarse
- Simple

Desired: *automatic*, *continuous* adaptation of the *intelligence* in a game (Game AI)

AI in Games

Game AI Levels

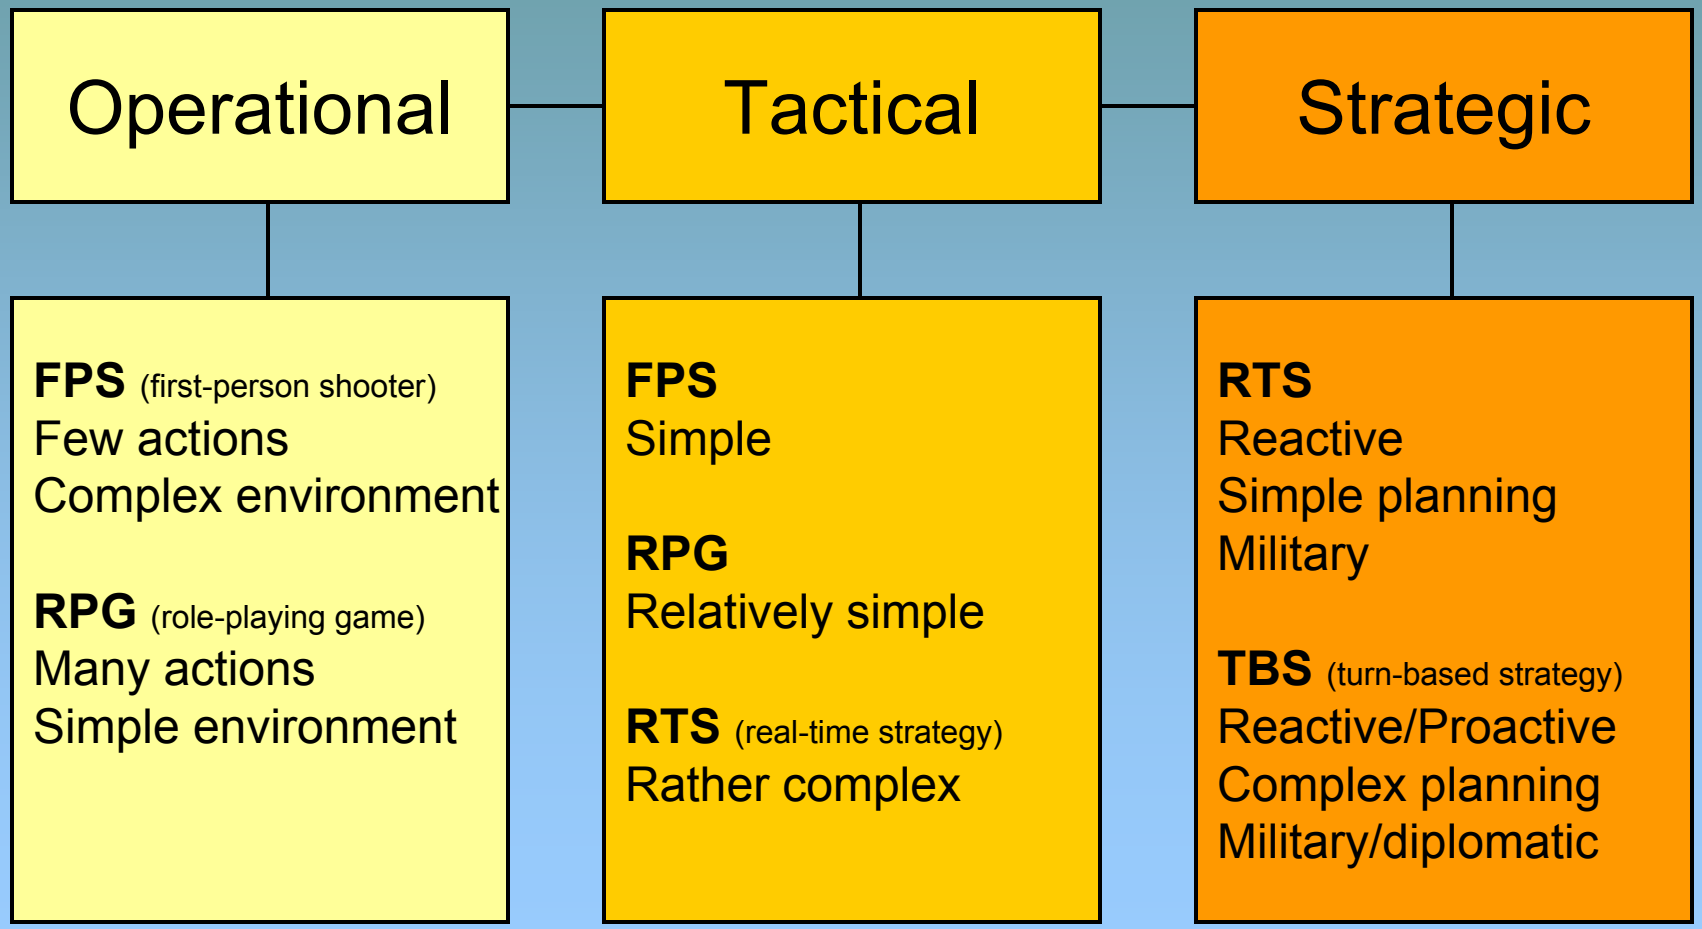

Adaptive AI

- Self-correction
  - Automatic repair of “exploits”
- Creativity
  - Being able to deal adequately with new situation
- Scalability
  - Appropriate challenge for both weak and strong human players

Should adapt to the human player, therefore, *during gameplay*

Adaptive AI Techniques

- Neural Networks
  - Few actions
  - Slow adaptation
- Evolutionary Algorithms
  - Generate many inferior solutions
  - Slow adaptation
- Reinforcement Learning
  - Based on rewarding good behaviour and punishing bad behaviour
  - Can learn during gameplay

Problem of Complexity

- Huge state-action space
- Non-deterministic
- Incomplete information
- Multiple parallel agents
- Real-time

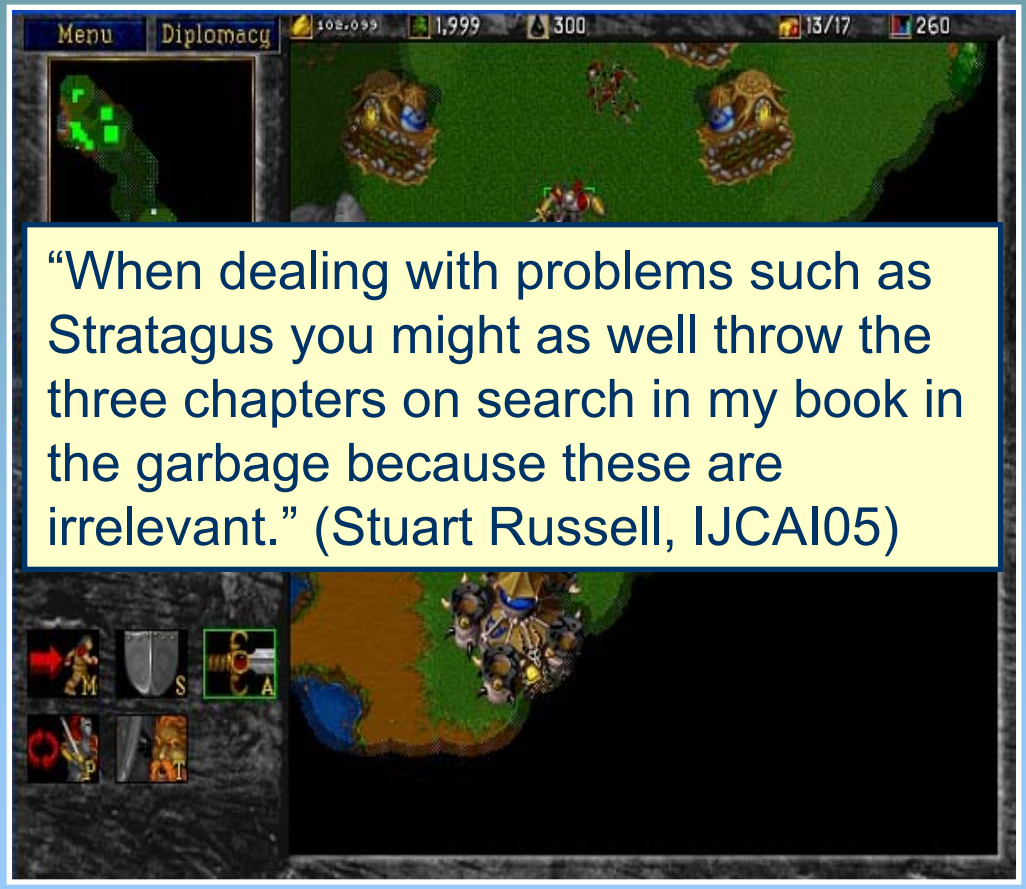

"When dealing with problems such as Stratagus you might as well throw the three chapters on search in my book in the garbage because these are irrelevant." (Stuart Russell, IJCAI05)

Dynamic Scripting

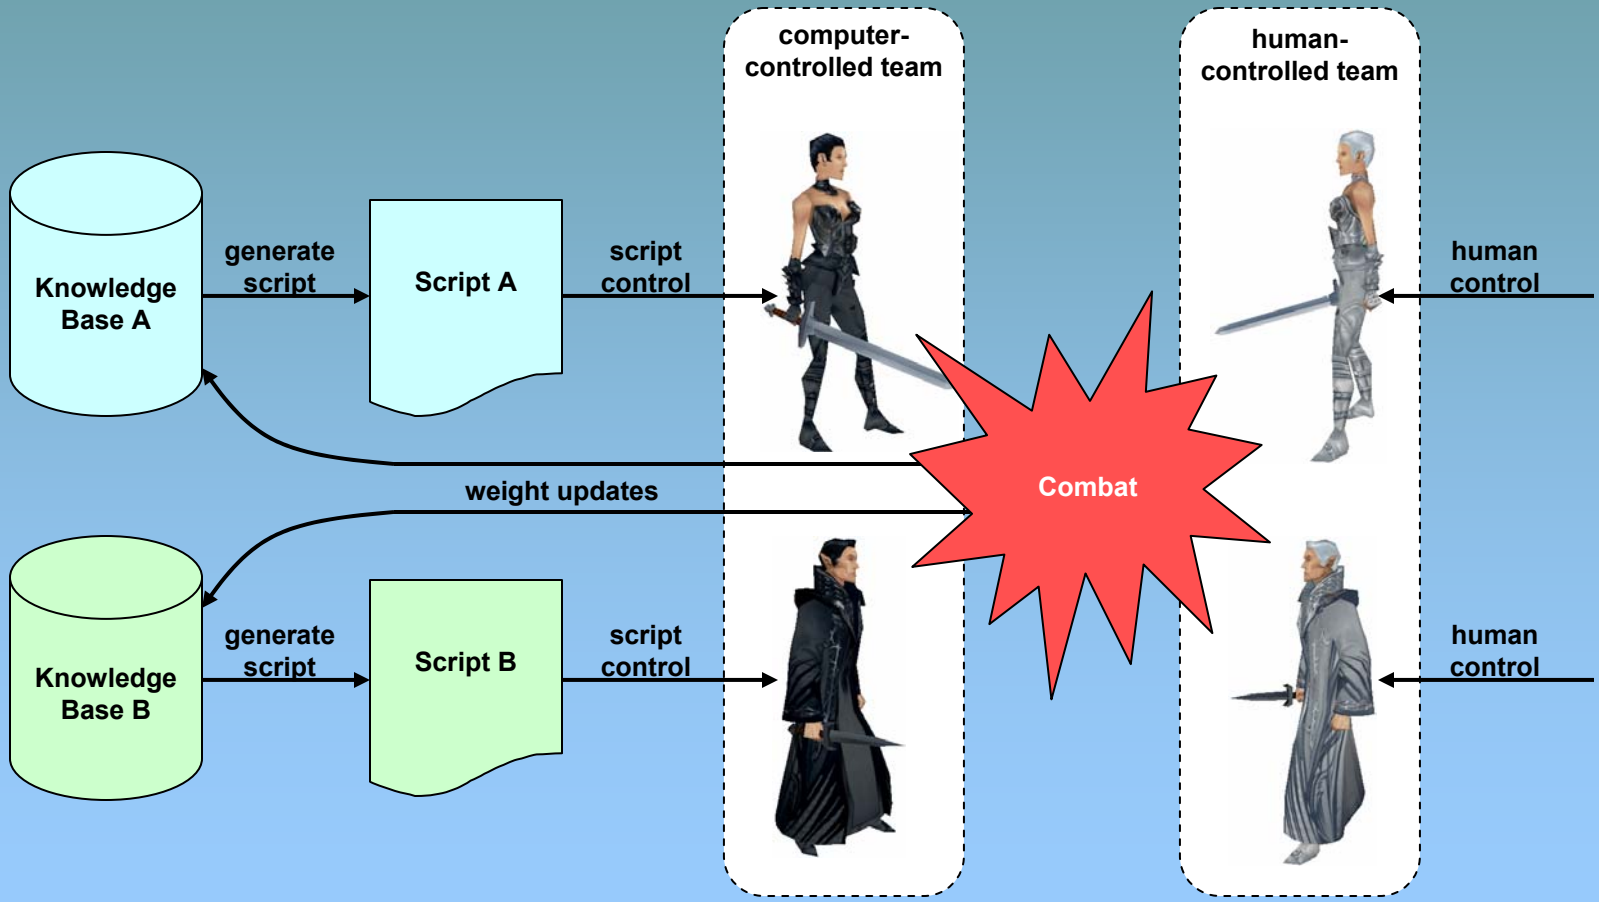

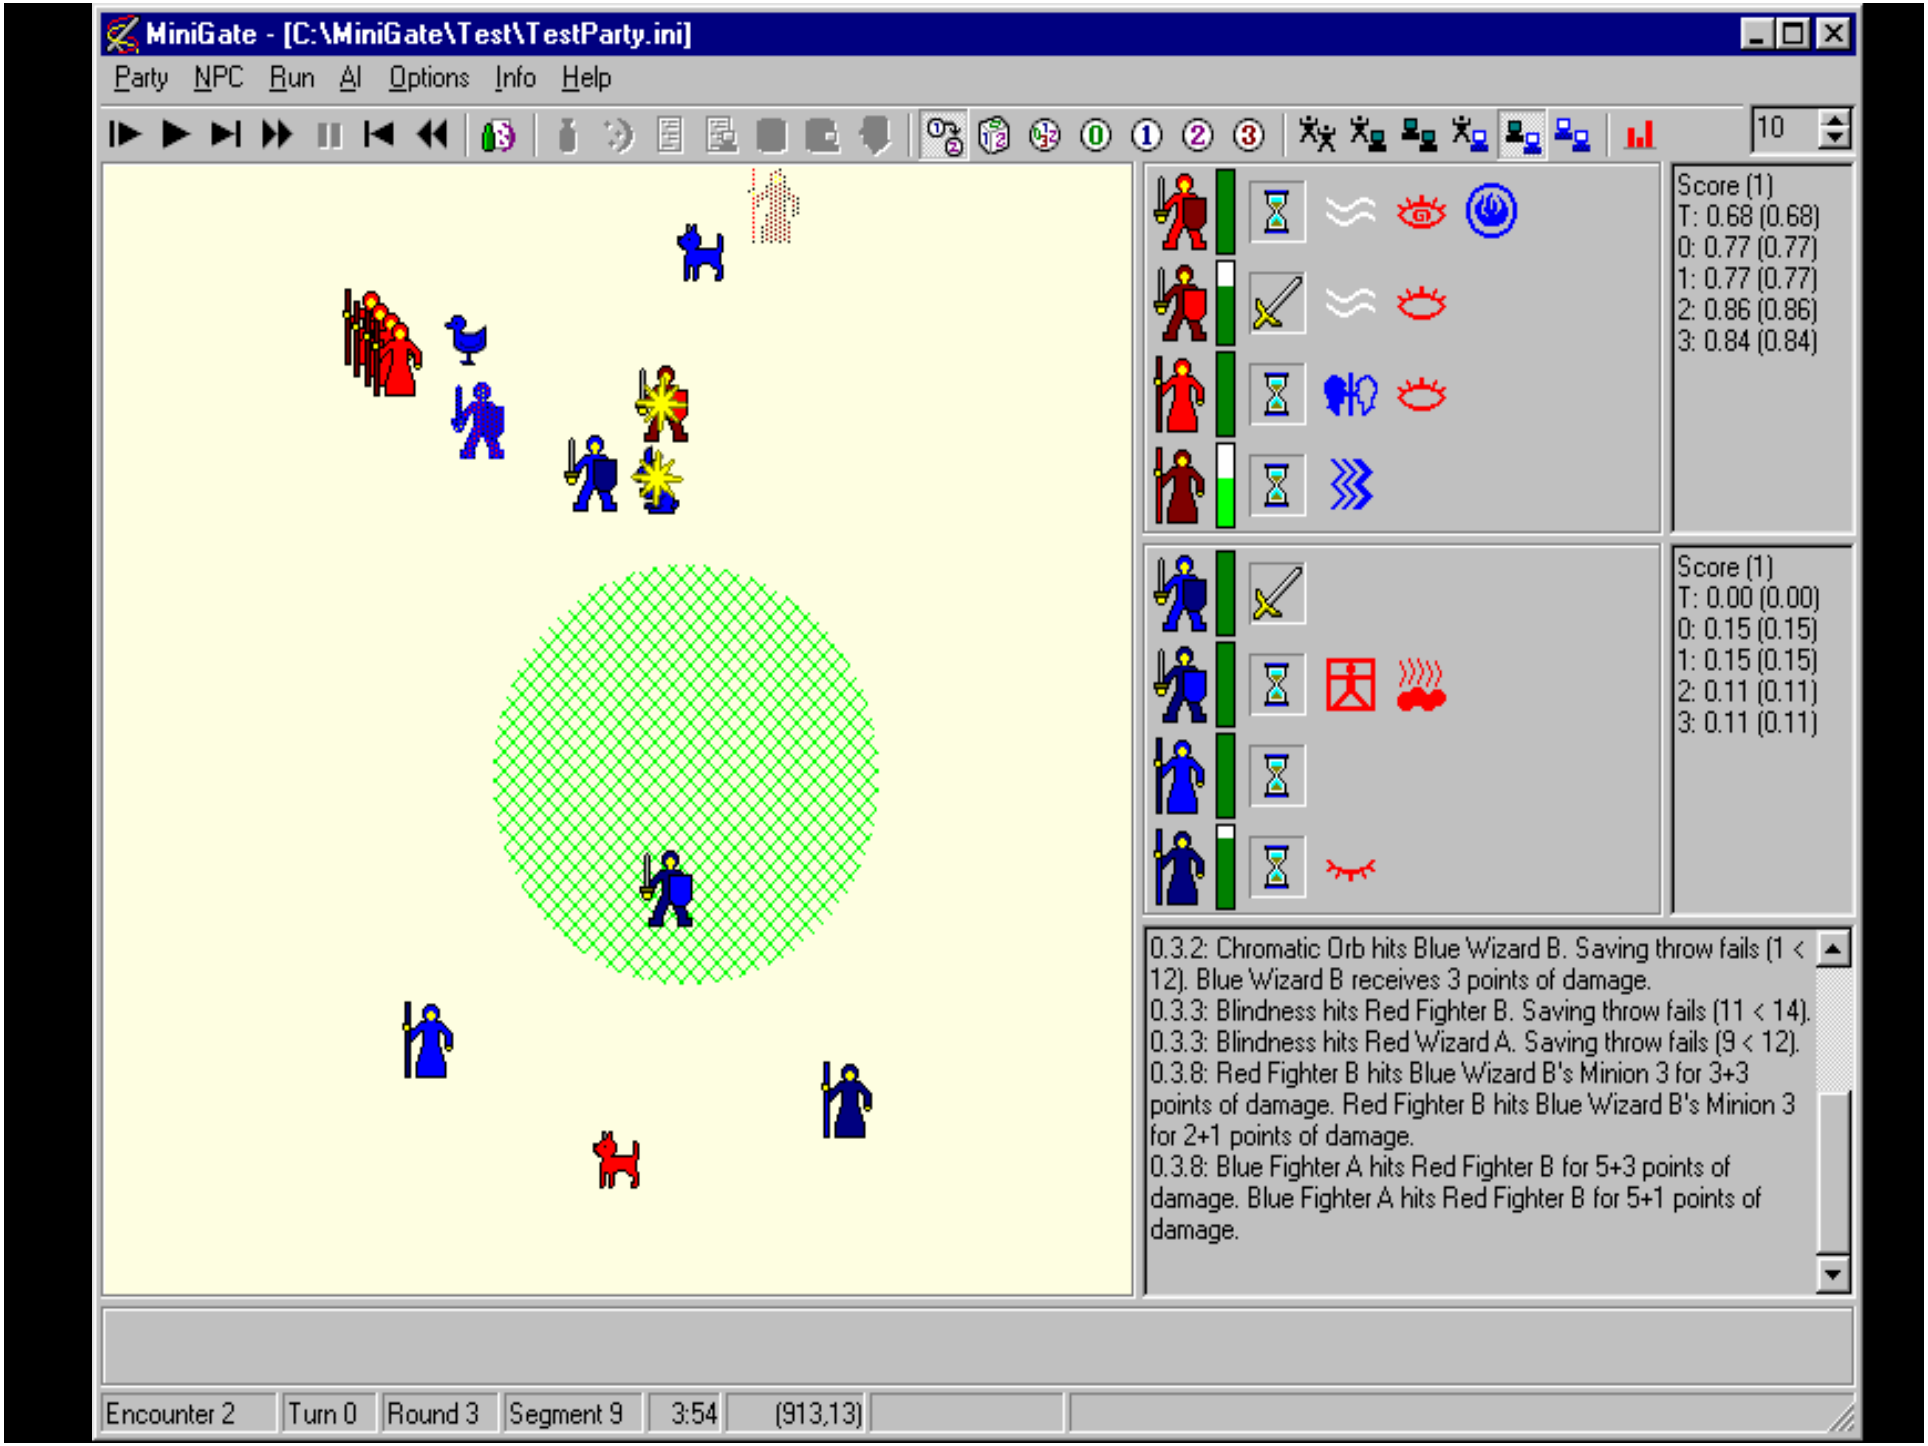
MiniGate - [C:\MiniGate\Test\TestParty.ini]
Party NPC Run AI Options Info Help
10
Score (1)
T: 0.68 (0.68)
0: 0.77 (0.77)
1: 0.77 (0.77)
2: 0.86 (0.86)
3: 0.84 (0.84)
Score (1)
T: 0.00 (0.00)
0: 0.15 (0.15)
1: 0.15 (0.15)
2: 0.11 (0.11)
3: 0.11 (0.11)
0.3.2: Chromatic Orb hits Blue Wizard B. Saving throw fails (1 < 12). Blue Wizard B receives 3 points of damage.
0.3.3: Blindness hits Red Fighter B. Saving throw fails (11 < 14).
0.3.3: Blindness hits Red Wizard A. Saving throw fails (9 < 12).
0.3.8: Red Fighter B hits Blue Wizard B's Minion 3 for 3+3 points of damage. Red Fighter B hits Blue Wizard B's Minion 3 for 2+1 points of damage.
0.3.8: Blue Fighter A hits Red Fighter B for 5+3 points of damage. Blue Fighter A hits Red Fighter B for 5+1 points of damage.
Encounter 2
Turn 0
Round 3
Segment 9
3:54
(913,13)

Simulation Results

Starting with all weights equal, the most challenging tactic is consistently defeated after 50 fights on average, with a quite low standard deviation

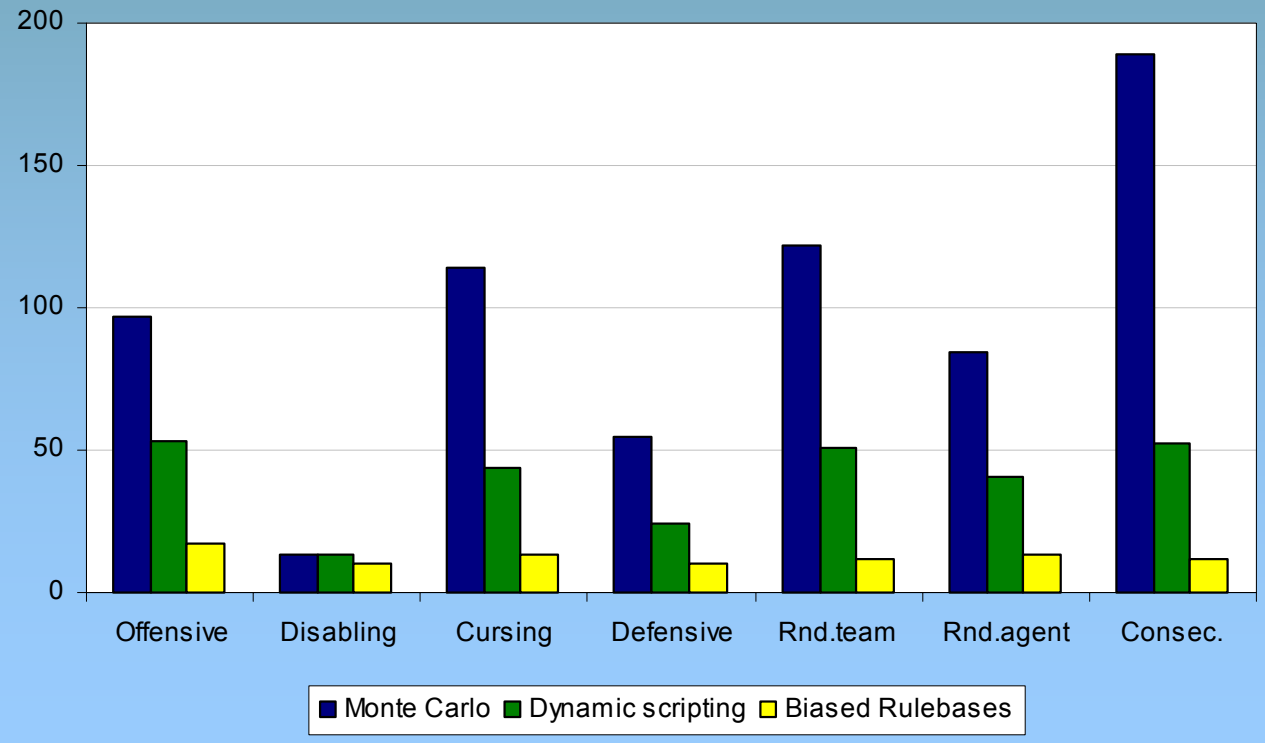


UNPAUSED
PAUSED
UNPAUSED
PAUSED


Map
Igor: Round 1, ladies and gentlemen!
PAUSED
Blanche the Cleric casting Summon Creature I

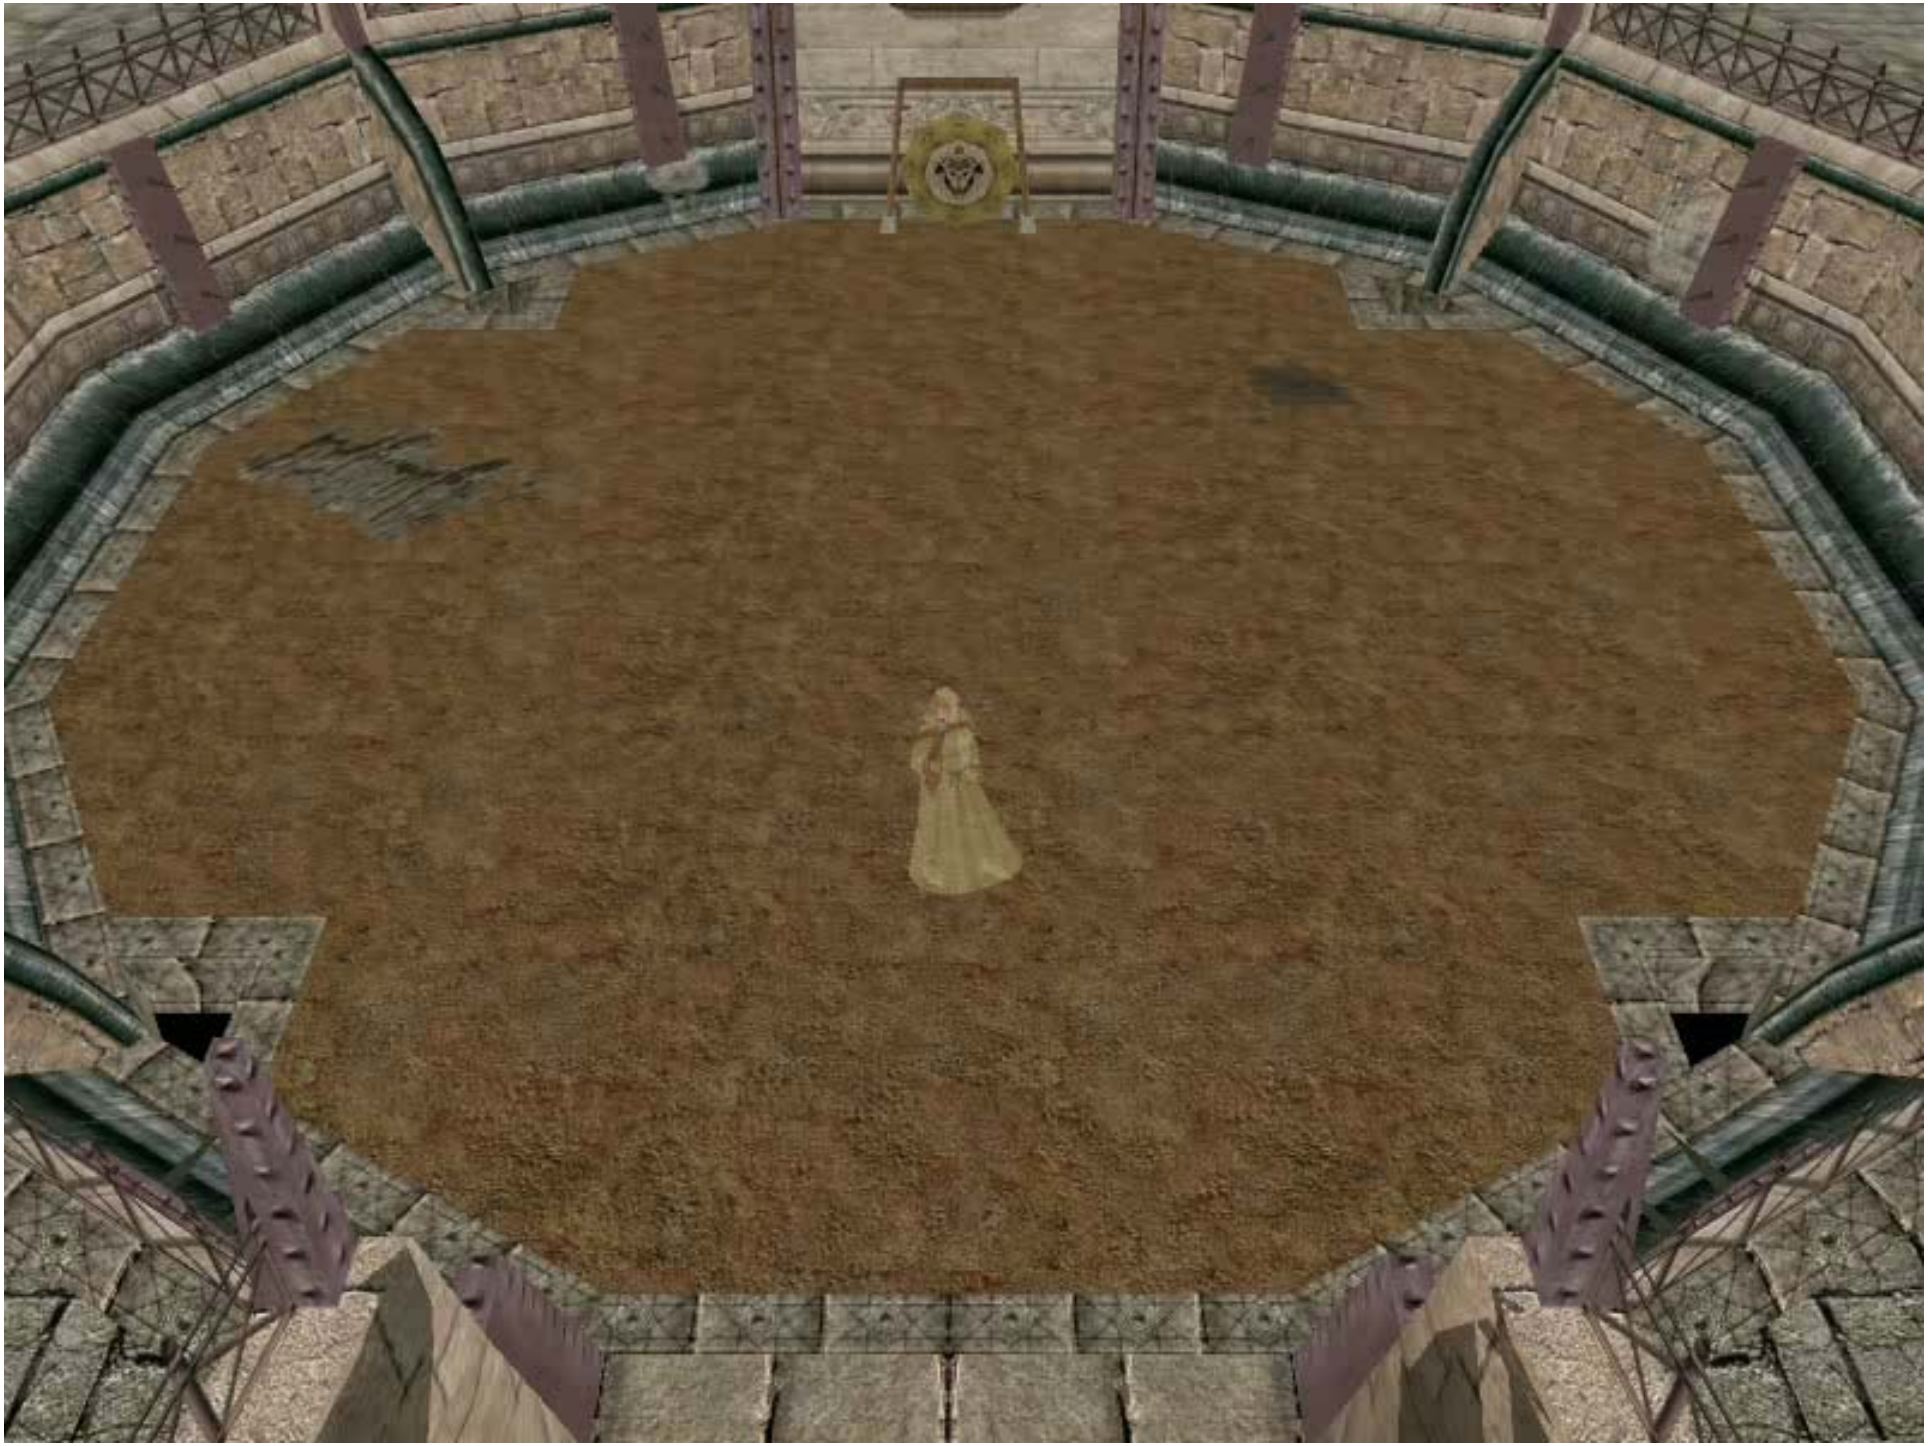

Automatic Scaling of Game AI

- “High-fitness penalising”
  - Award the highest fitness to the “most equal” AI, instead of to the “best” AI
- “Weight clipping”
  - Increase AI variety when the computer plays too well
  - Decrease AI variety when the computer plays badly
- “Top culling”
  - Remove the currently “best” knowledge when the AI plays too well
  - Reactivate the “best” knowledge when the AI plays badly

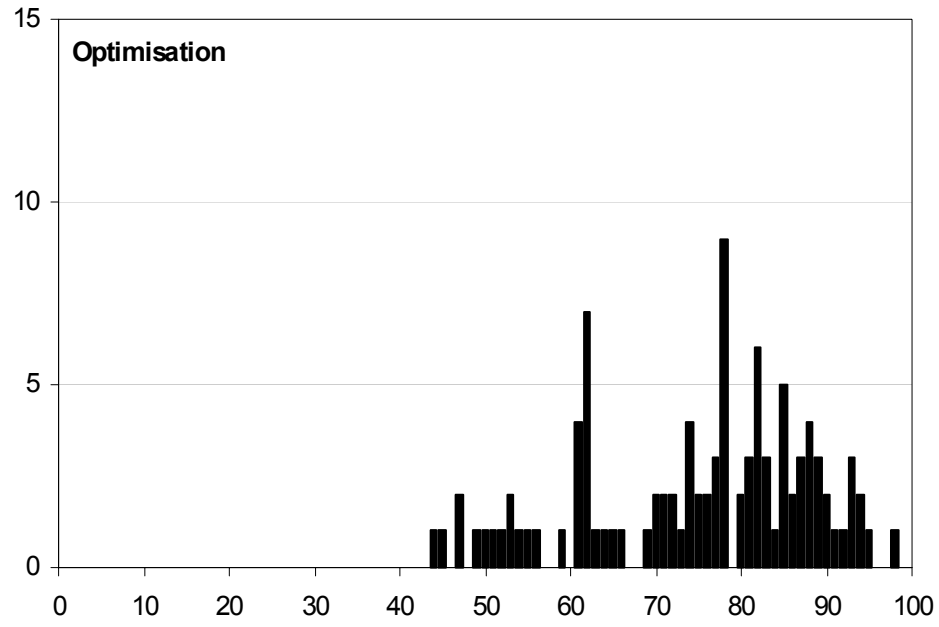
Optimisation
15
10
5
0
0 10 20 30 40 50 60 70 80 90 100

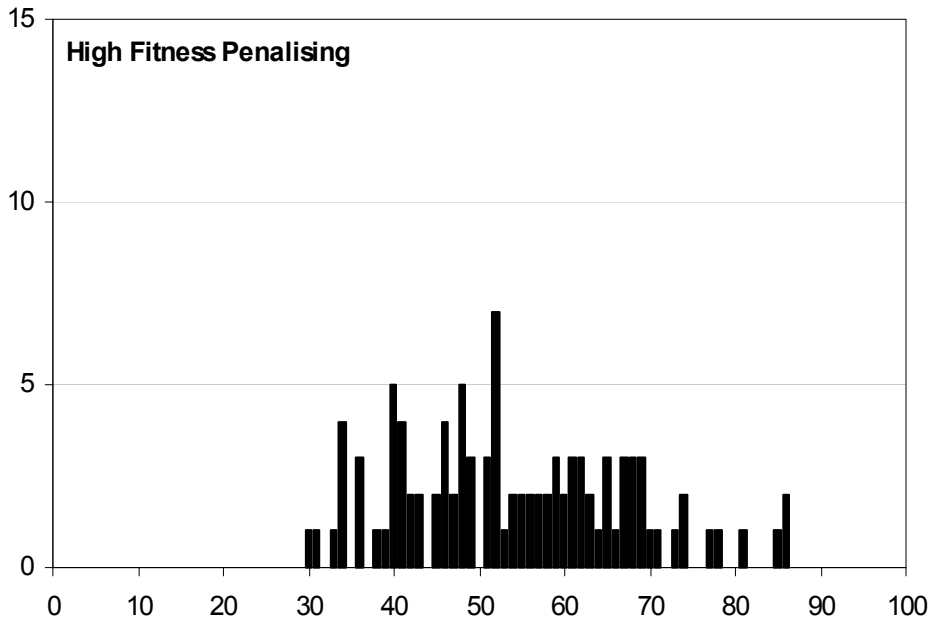
High Fitness Penalising
15
10
5
0
0 10 20 30 40 50 60 70 80 90 100

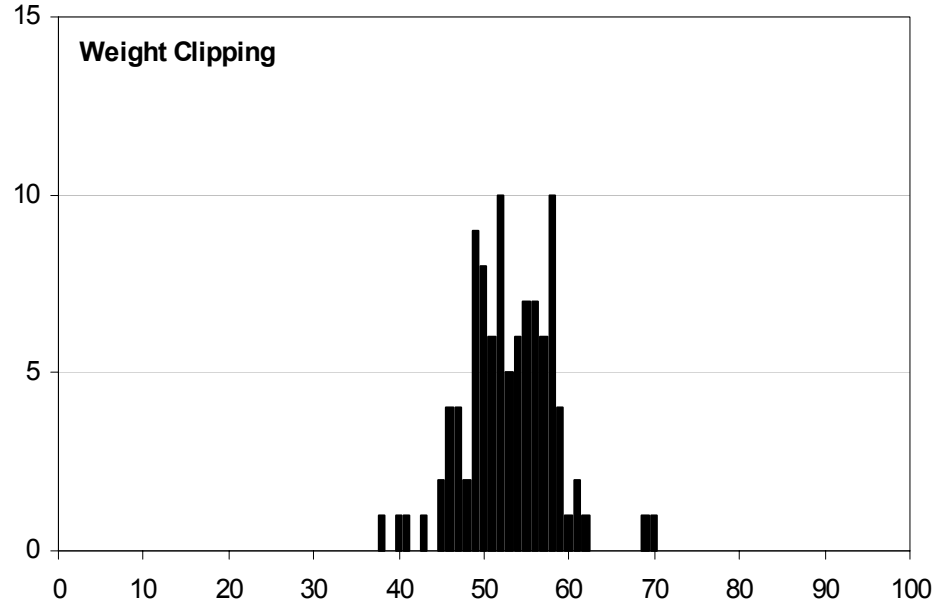
Weight Clipping
15
10
5
0
0 10 20 30 40 50 60 70 80 90 100

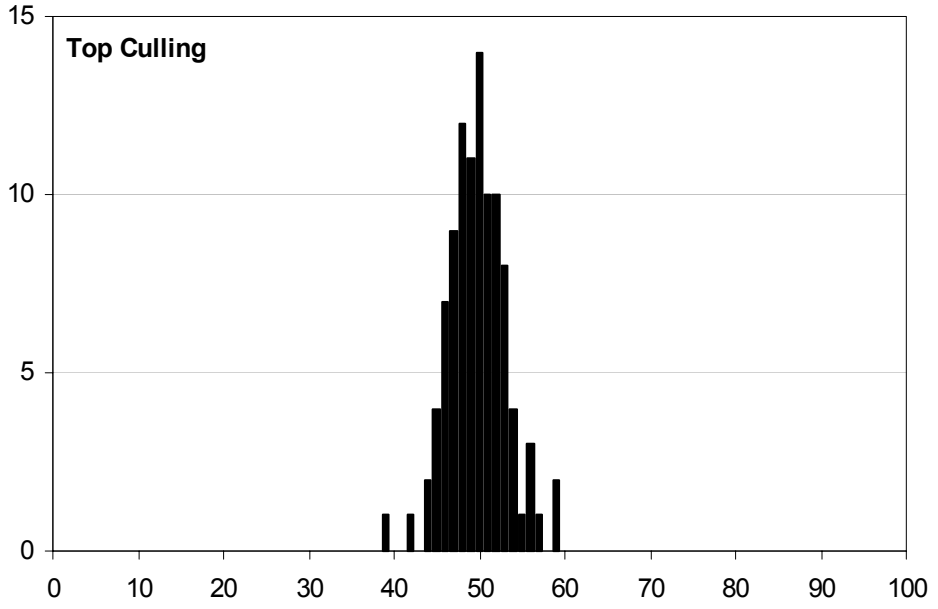
Top Culling
15
10
5
0
0 10 20 30 40 50 60 70 80 90 100

Resultaten Schaling

- *Without automatic scaling*, dynamic scripting wins against all tactics
- With *high-fitness penalising* an even game is achieved against 2 out of 8 tactics
- With *weight clipping* an even game is achieved against 7 out of 8 tactics
- With *top culling* an even game is achieved against 8 out of 8 tactics, combined with the lowest standard deviation

Use of Automatic Scaling

- Not useful against strong players
- Players might lose a sense of accomplishment when they find out it is activated
- Best results against novices

- For best teaching results, should the AI play just a bit stronger than the human player?
  - If so, how much stronger?
  - Is this equal for all human players?

Tactical and Strategic

- Adaptive AI on a Tactical Level
  - Team configuration in Quake III with symbiotic AI
- Adaptive AI on a Strategic Level
  - Planning in Stratagus with dynamic scripting
  - Automatic design of dynamic scripting knowledge bases for Stratagus using evolutionary algorithms

Architecture for Adaptive Game AI

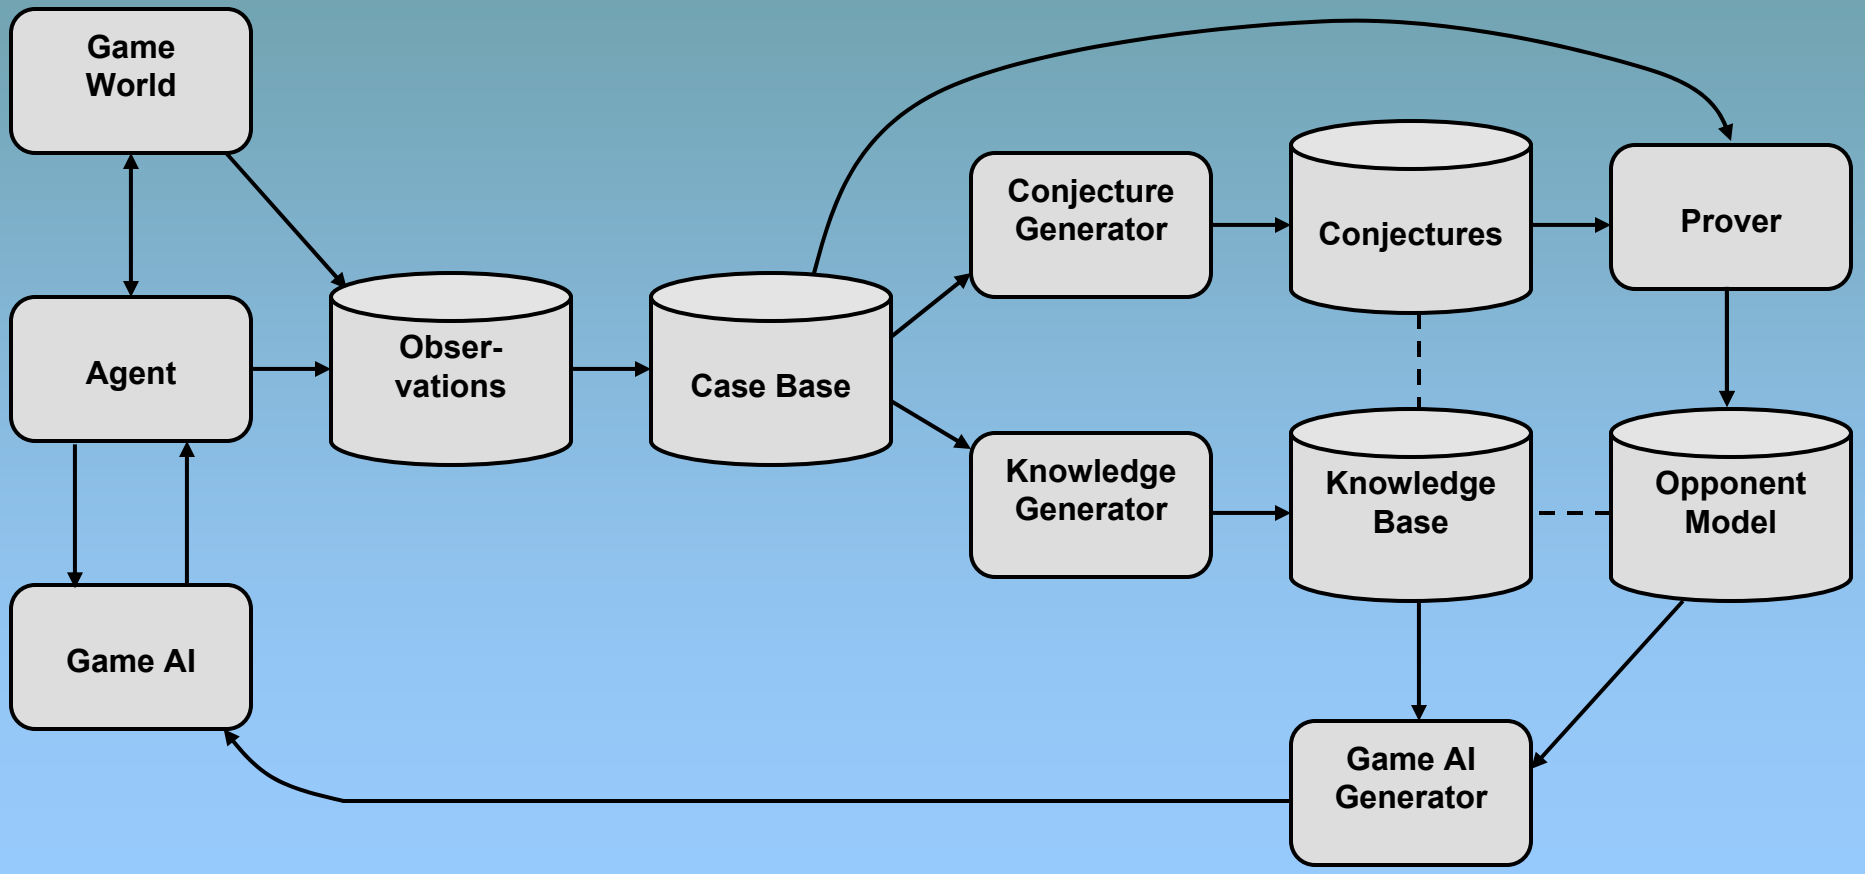

Conclusions

- Enjoyment of a game is increased if the challenge level is appropriate for the human who plays the game
- To let a human player learn from a game, or to achieve maximum enjoyment from a game, the game should adapt to the specific human who plays it
- Complexity reduction is necessary to achieve such
- Specific forms of reinforcement learning are suitable for this purpose (e.g., dynamic scripting)