Artificial Intelligence and Games
Introduction

Georgios N. Yannakakis
@yannakakis

Julian Togelius
@togelius

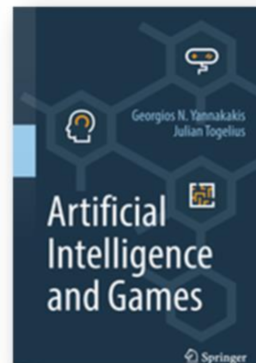

These are the slides accompanying the book Artificial Intelligence and Games through the gameaibook.org website

Why a course on AI *in* and *for* Games?

Artificial Intelligence is at the very core of modern game design and development. It plays and tests games, it supports the design of games and it analyses the way we play them. AI techniques used for these tasks include **expert domain-knowledge systems, search and optimization, computational intelligence, and machine learning**. This course aims to introduce students to the theory and practice of **game artificial intelligence.**

Some reasons why a course on game AI is time-relevant and important

Learning Objectives

- Describe and theorize on the AI algorithms covered in class.
- Identify tasks that can be tackled through AI techniques and select the appropriate technique for the problem under investigation.
- Compare the performance of different AI techniques and reflect on their suitability for game AI development.
- Design and implement efficient and robust game AI algorithms.

Some potential learning objectives for a graduate course on game AI

Class Core Aim

- Give a comprehensive overview of **AI for games** and **games for AI**, as well as hands-on experiences with new technology
- Enable new exciting inventions and discoveries

A potential core aim for a course on artificial intelligence and games

Let us start by discussing AI first…

Artificial Intelligence

Making **computers** able to do things which currently only **humans** can do

A popular definition of AI

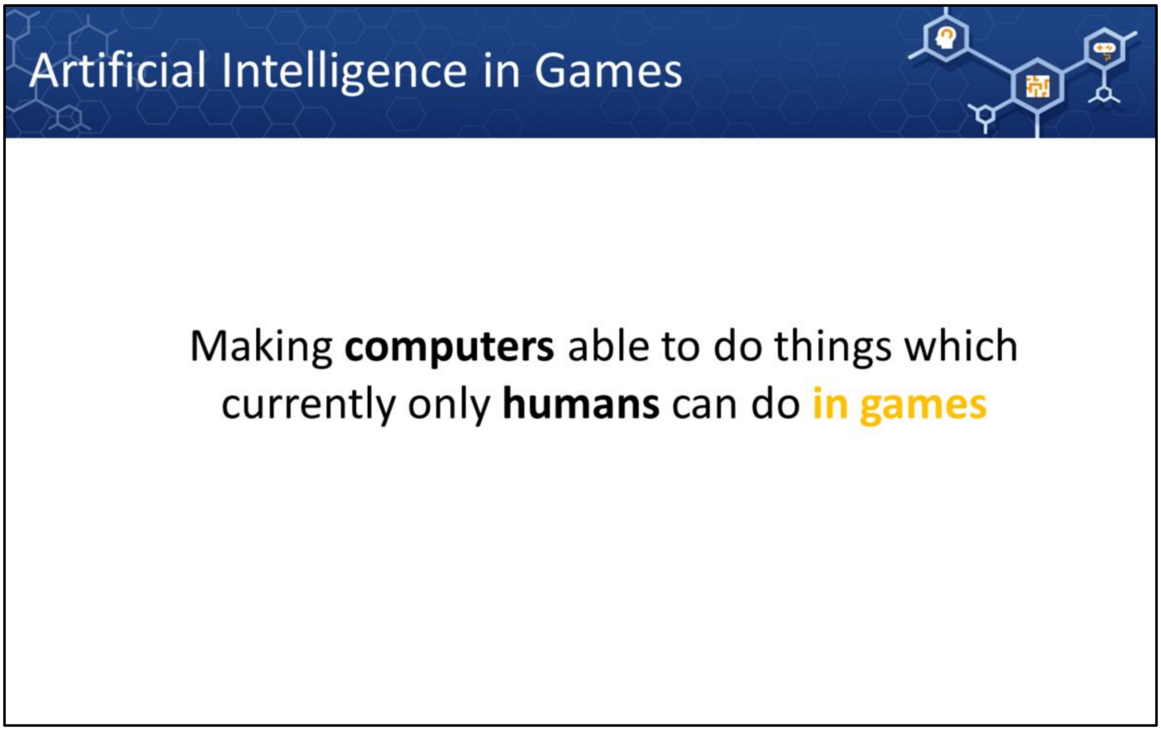

AI as employed to games

What do Humans Do with Games?

- Play them
- Study them
- Build content for them
  - levels, maps, art, characters, missions…
- Design and develop them
- Do marketing
- Make a statement
- Make money!

Here is a non-inclusive list of things humans can do with games. What if AI could take these roles?

[see Section 1.4 for more details]

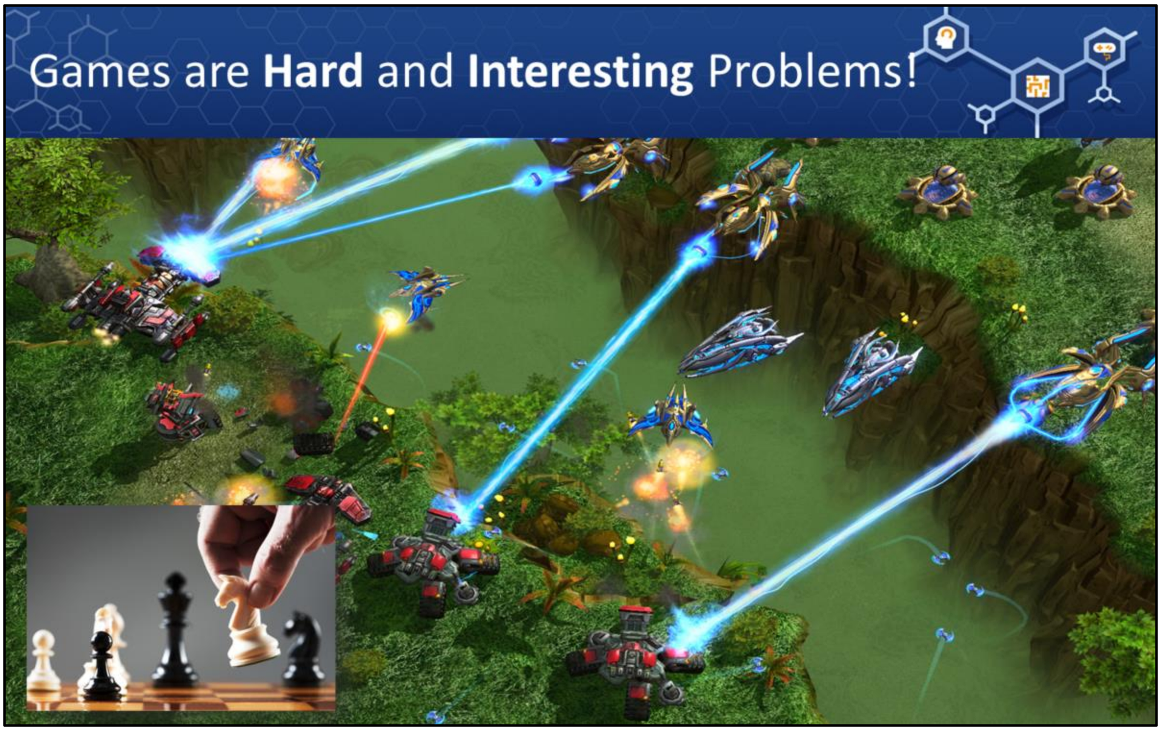

[see Section 1.3.1 for more details]

Just the AI playing most digital games is beyond a NP-hard complexity

- Hard evaluation functions (goodness movement, interestingness, engagement)
- Large search spaces
- Real-time!
- Made to challenge human cognitive/social/emotional abilities

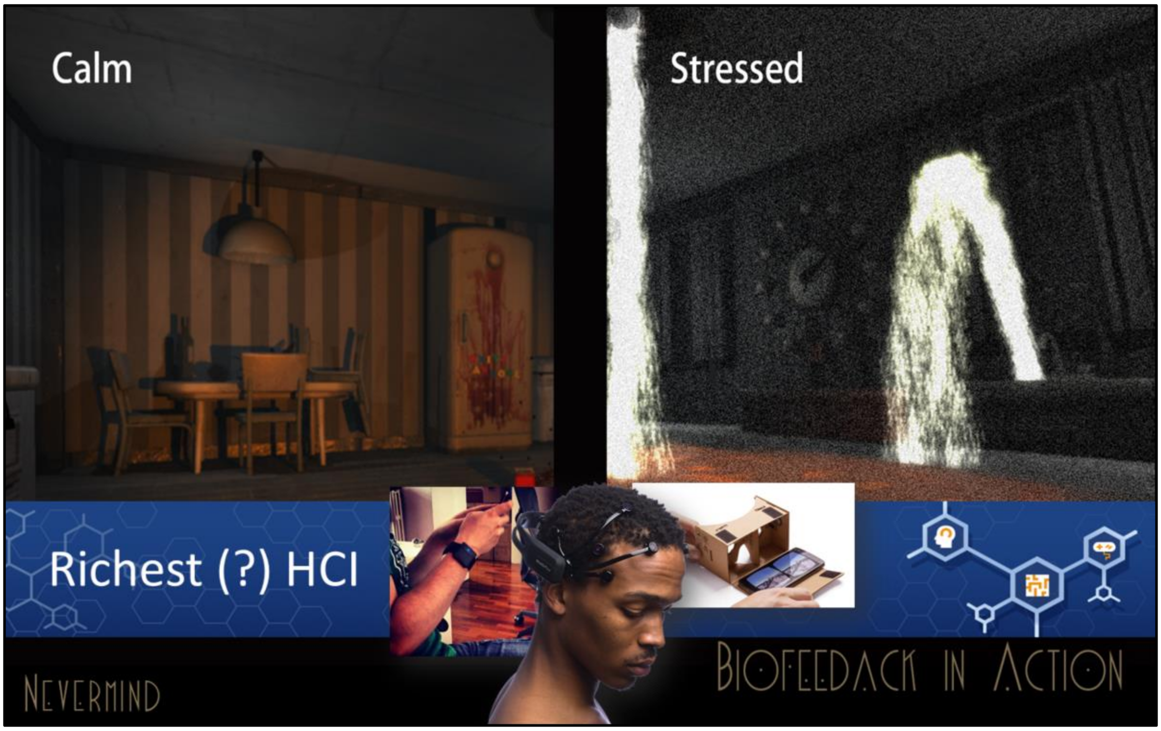

[see Section 1.3.2 for more details]

Games : dynamic media by definition and, arguably, offer one of the richest forms of human-computer interaction (HCI)

The richness of interaction is defined in terms of

- the available options a player has at any given moment (i.e. game and action space);
- the ways (modalities) a player can interact with the medium - keyboard, mouse and tablet-like haptics, to game controllers, physiology such as heart rate variability, body movement such as body stance and gestures, text, speech […]

Games offer one of the best and most meaningful domains for the realization of the affective loop

Image: Nevermind biofeedback game: http://nevermindgame.com/biofeedback/

[see Section 1.3.2 for more details]

Games are now **accessible** and **fashionable**.

- Games are “sexy” for the general populace.
- Any smartphone these days will have many games installed.
- People love discussing their game experiences at the office, at school, at the bar.
- They are proud to wear t-shirts displaying their favorite game.
- Even museums consider them a valid art form as is the case of the Smithsonian Art Museum which displayed The Art of Video Games in 2012.
- Serious Games for Education/Health/Training are getting increasingly important and accessible

[see Section 1.3.3.1 for more details]

A bit more on content:

- Content costs: money and effort !
- Open boundaries for creativity as content for each creative domain comes in different representations, under different sets of constraints and often created in massive amounts

Bottom right image: No Man's sky : over 18 quintillion(1.8×10^{19}) planets, many with their own set of flora and fauna

Games and AI Areas

Perception
Signal Processing
Machine Learning
Artificial Psychology
Planning and Search
KR and Reasoning
Natural language processing
Navigation
...

[see Section 1.3.4 for more details]

Unlike some more narrow benchmarks, games challenge all core areas of AI. This can be seen by taking a number of broadly accepted areas of AI and discussing the challenges available for those areas in games.

All areas are **advanced** via games…

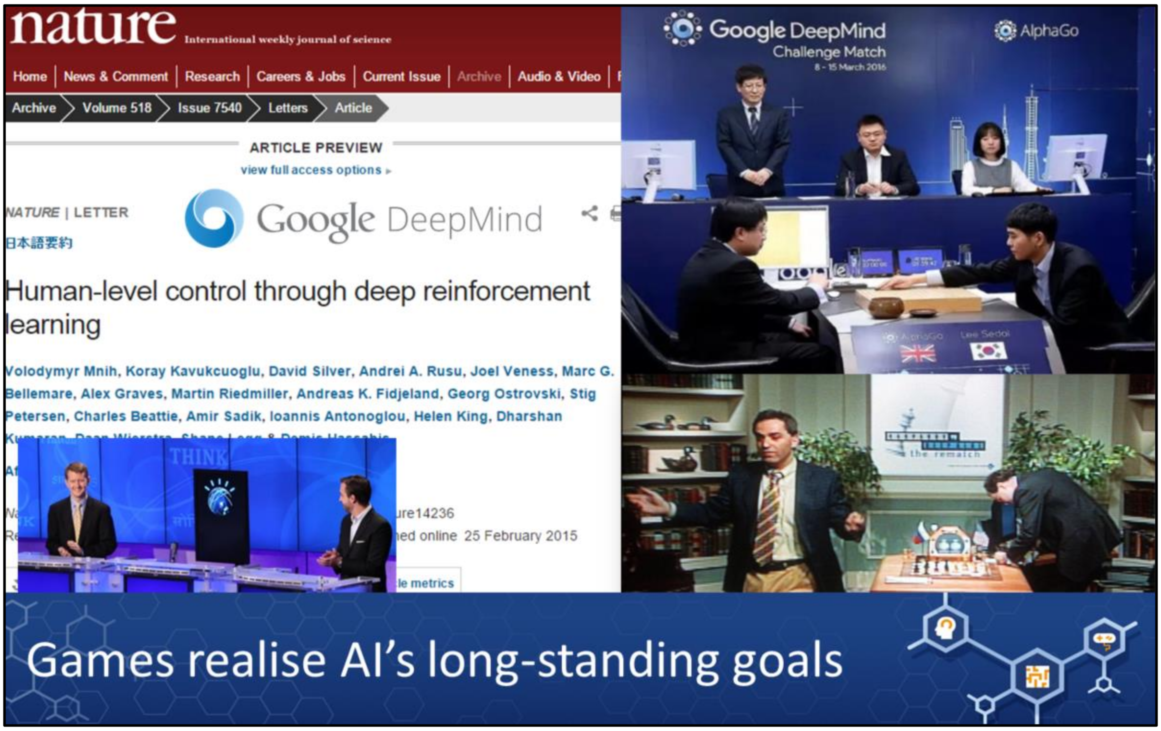

[see Section 1.3.5 for more details]

Social and Emotional intelligence [Section 1.3.5.1]

Computational creativity [Section 1.3.5.2]

General intelligence [Section 1.3.5.3]

All AI milestones are based on or derived from games (Deep Blue, Kinect , Jeopardy, Chinook-Checkers, AlphaGo, …)

[see Section 1.3.5.2 for more details]

On computational creativity in games:

- Games is the domain where
- Art meets science and play …
- Problem solving meets creativity
- The outcome is experienced (via rich interaction)!

In summary: An AI that can play almost any video game and create a wide variety of video games is, by any reasonable standard, intelligent. That task can be seen as AI complete

Artificial Intelligence **for** Games

[see Section 1.4 for more details]

[see Section 1.4.1 for more details]

AI can improve games in several ways:

- When the AI adds to the commercial value of the game, it contributes to better game reviews, and it enhances the experience of the player.
- An unconventional and effective solution to an NPC task can often be a critical factor that shapes management, marketing and monetization strategies

AI roles and tasks:

- AI plays games with two core objectives in mind: **play well** and/or **play believably** (or human-like, or interestingly)
- AI can control either the player character or the NPC
- AI that plays well: **empower dynamic difficulty adjustment** and **automatic game balancing**
- AI that plays believably or human-like: **player experience debugging** or **demonstration of realistic play**

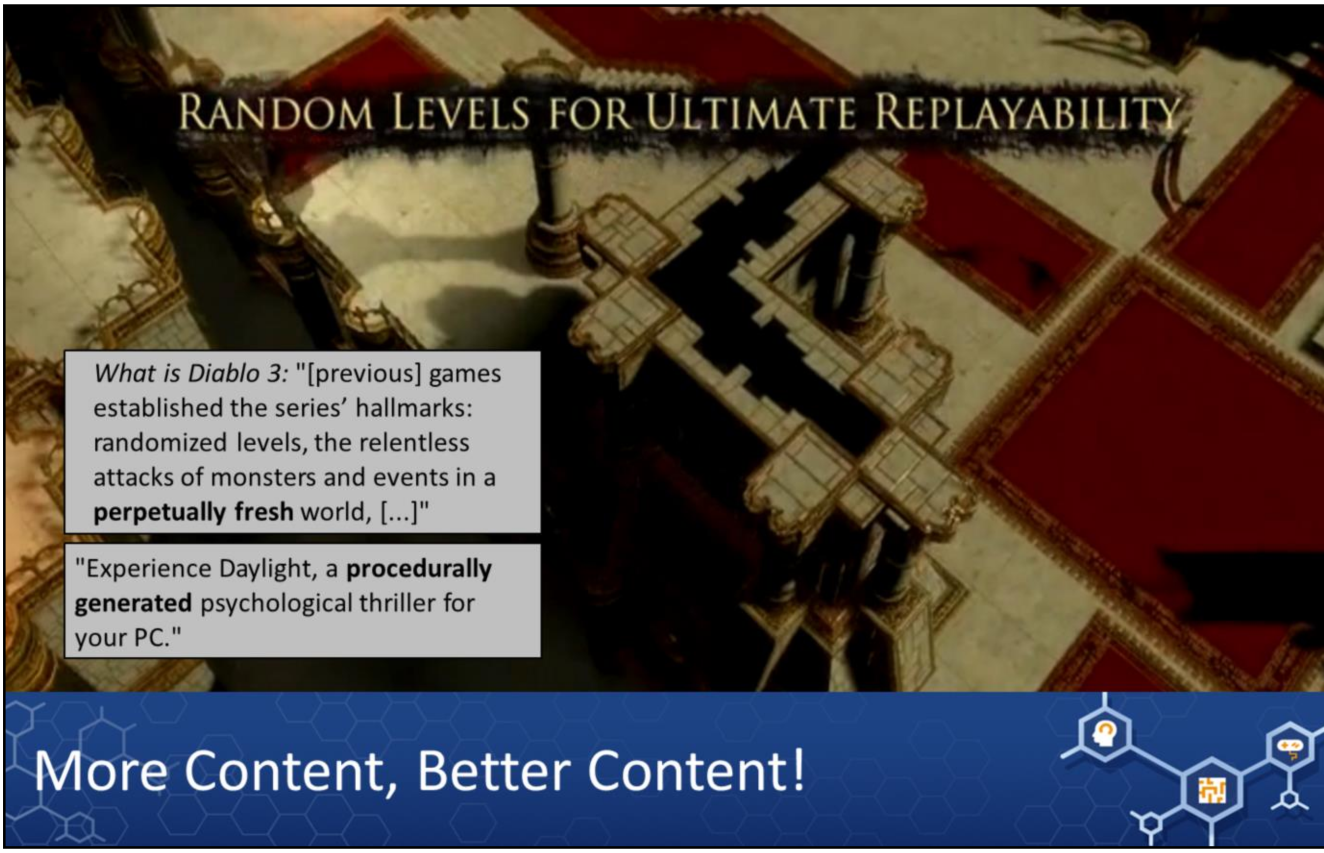

[see section 1.4.2 for more details]

- Games are unique → they **invented** procedural content generation (PCG)
- PCG is a **commercial necessity:** highly competitive marketplace (fast development cycles), replayability, retention (online games)…
- A constant stream of new content is needed!
- The game industry **proudly** displays its AI.

Some reasons for game designers and developers to be interested in AI, and PCG in particular (see more in Chapter 4):

- **Memory consumption** (as e.g. in *Elite)*
- Content generation might foster or further inspire **human creativity**

[see Section 1.4.3 for more details]

- The holy grail of game design—player experience—can be improved and tailored to each player
- The use of AI for the understanding of player experience can drive and enhance the design process of games […]
- Data derived from games provide a new and complementary way of designing games, of making managerial and marketing decisions about games, of affecting the game production, and of offering a better customer service […].
- AI-informed decisions are based on **evidence** rather than *intuition* (via game analytics and game data mining) for better design, development and QA procedures.
- Game development is boosted and improved (as a whole).
- In summary: AI-enabled and data-driven game design can directly contribute to **better games**.

[See section 1.2 and 1.2.1 for more details]

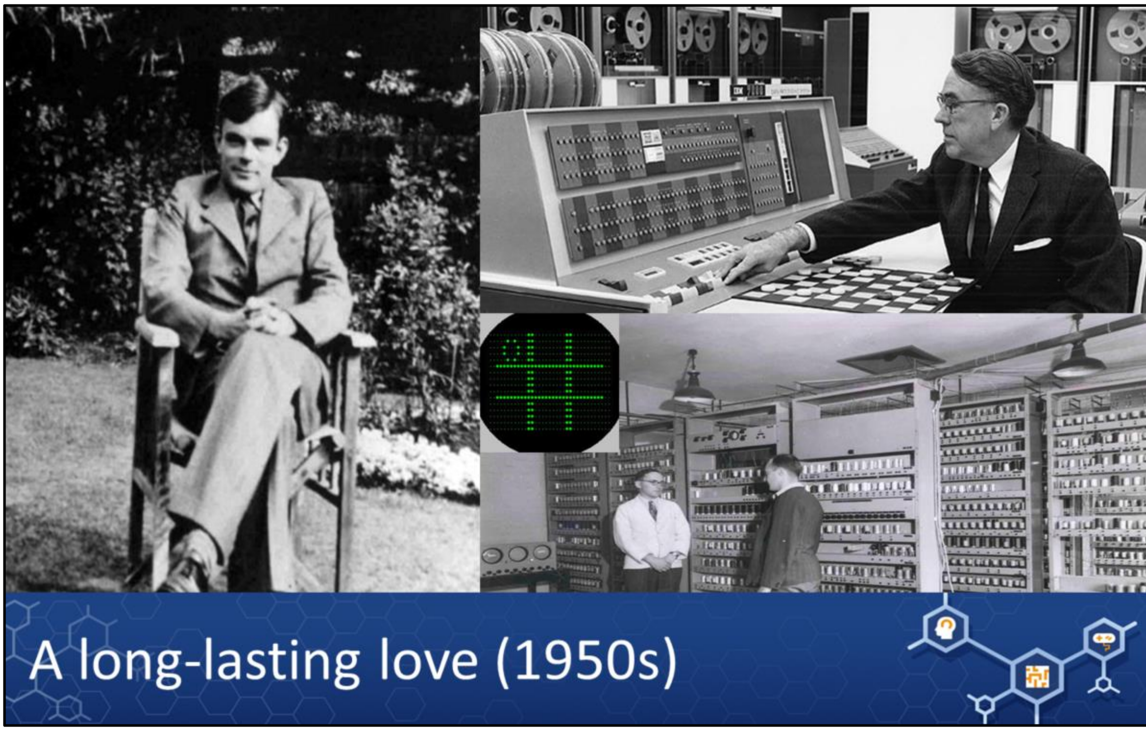

[see Section 1.2.1 for more details]

- 1951: Alan Turing, (re)invented the Minimax algorithm and used it to play Chess (TUROCHAMP program – one move chess analyser) [image: left]
- 1952: The first software that managed to master a game was programmed by **Alexander "Sandy" Douglas** on a digital version of the Tic-Tac-Toe game and as part of his doctoral dissertation at Cambridge. [image: bottom right]
- 1959: Arthur Samuel was the first to invent the form of machine learning that is now called reinforcement learning using a program that learned to play Checkers [image: top right]

[see Section 1.2.1.1. for more details]

- 1992: TD-Gammon employs an ANN trained via TD learning by playing backgammon against itself a few million times. TDGammon [image: bottom right] managed to play at a level of a top human backgammon player
- 1994: Over three decades of research on tree search the **Chinook** Checkers player managed to beat the World Checkers Champion **Marion Tinsley -** Checkers was eventually solved in 2007 [image: top left]
- 1997: IBM's **Deep Blue** (version of Minimax) beats Gary Kasparov. In the 44th move of game 1 – Kasparov realises he is loosing by a superior Chess player - the move was a result of a bug!) [image: bottom left]
- 2017: AlphaGo (deep reinforcement learning) won a three-game Go match against the world's number 1 ranking player **Ke Jie.** Go was the last great classic board game where computers have reached super-human performance. [image: top right]

The Digital Era - A milestone

- The Game Turing test was passed by two AI-controlled bot entries in *Unreal Tournament 2004,* in 2012!
- Neuroevolution was one of the approaches

[see Section 1.2.1.2 for more details]

2K bot prize (Unreal Tournament 2004)

Images: from the IEEE CIG 2008 conference competition [left: players; right: judges]

[see Section 1.2.1.2 for more details]

Notable milestones (playing games)

2014: Google DeepMind learned to play several games from the classic Atari 2600 video game console on a super-human

skill level just from the raw pixel input

2017: Ms Pac-Man is practically solved by the Microsoft Maluuba team using a hybrid reward architecture for RL

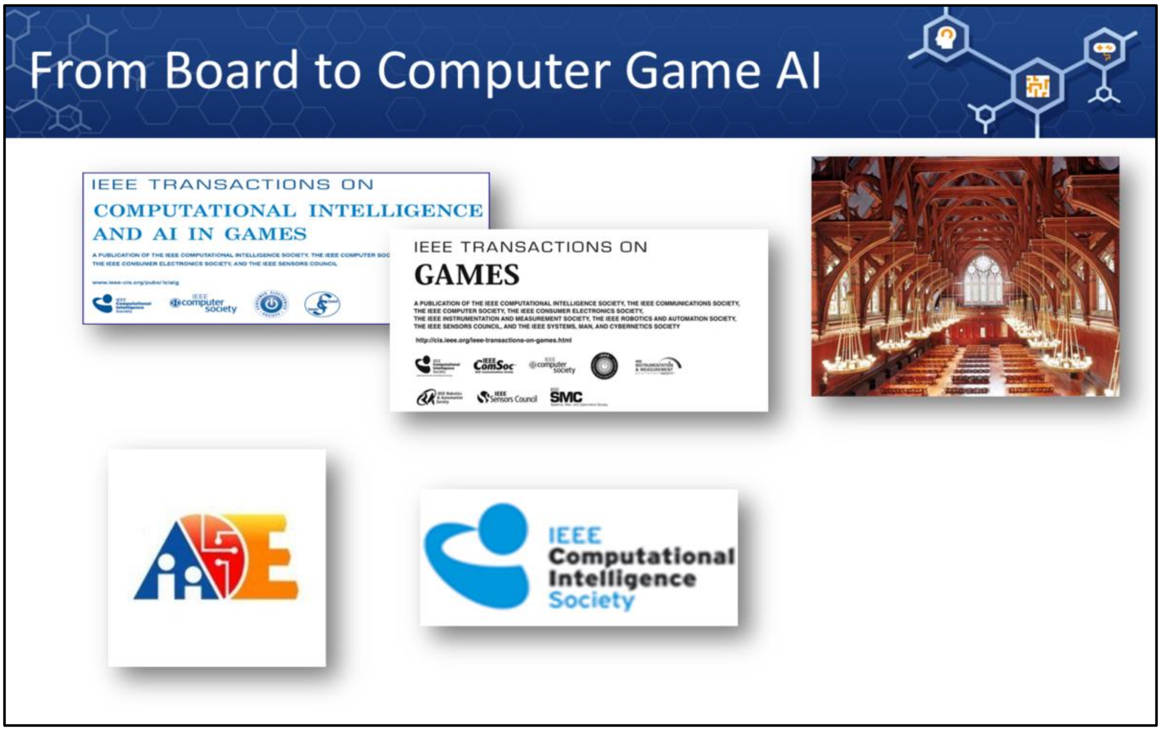

1983: The first video game conference occurred at Harvard's Graduate School of Education

2001: “Birth” of game AI. The seminal article by Laird and van Lent, “Human-level AI's killer application: Interactive computer games”

2000-2004: Early days: playing games, agent architectures for non player character (NPC) behaviour - sometimes within interactive drama- and pathfinding

2005: Birth of IEEE CIG and AIIDE conferences

2009: Launch of IEEE Transactions on Computational Intelligence and AI in Games

2018: Launch of IEEE Transactions on Games

IEEE CIG and AIIDE Competitions

A number of successful academic game AI competitions that run (and/or still running) at IEEE CIG and AIIDE conferences

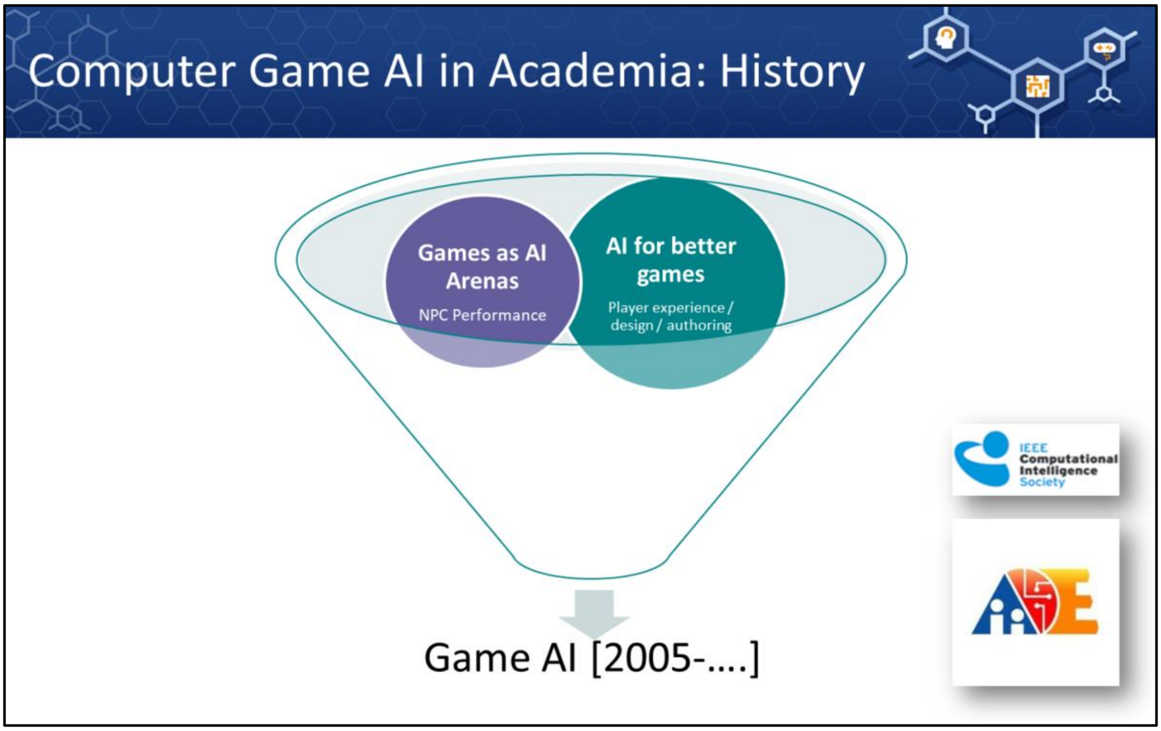

Brief history of Game AI roles/perspectives

2000-2004: Game AI = Non-Player Character (NPC) AI

2005-2012: A survey on Game AI papers published in IEEE CIG/AIIDE (premier game AI venues) identifies two main research tracks:

- Games as AI Arenas - 54% of papers for NPC control/pathfinding/decision making; Focus on NPC performance
- AI for better games - 46% of papers for non-traditional uses of AI; Focus on player experience/design/authoring

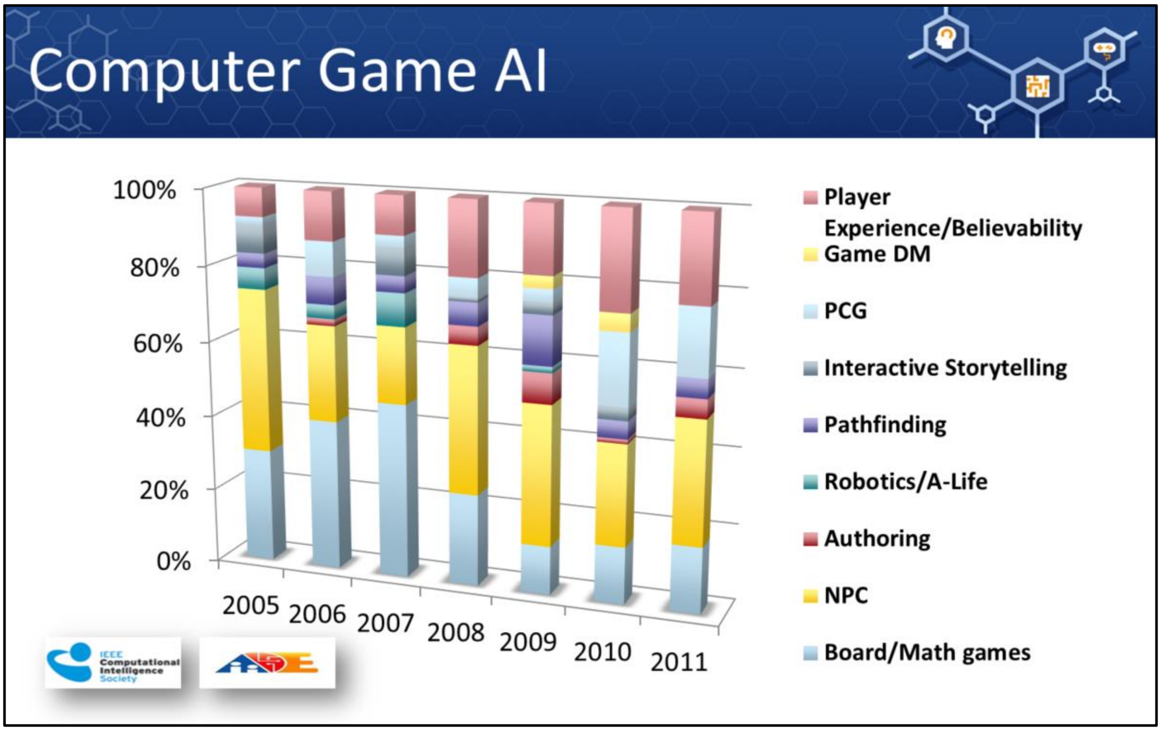

Graph: Percentage of papers published in AIIDE and IEEE CIG conference (the two premier conferences in game AI) throughout the years and across the different areas.

The field becomes more diverse as the years go by. A newer study (2012-present) would likely reveal a higher degree of diversity and maturity as a field.

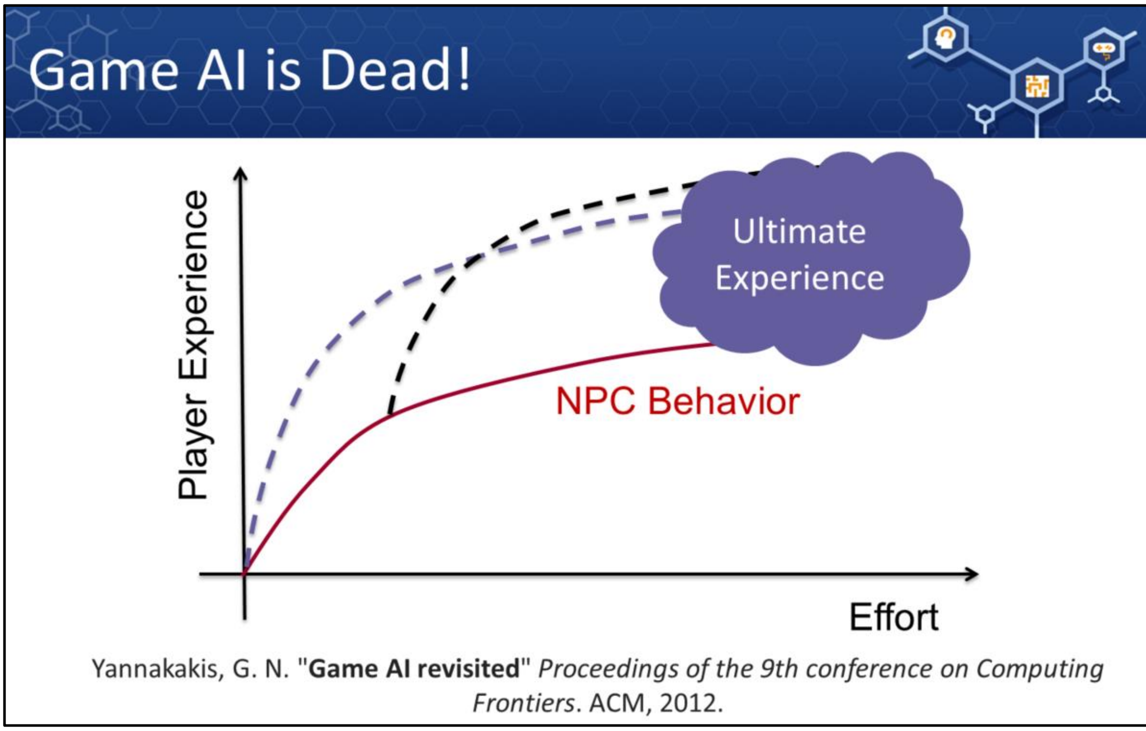

Graph explaining that while NPC behaviour can help us reach great levels of the ideal player experience there might be better ways (e.g. through procedural content generation, player modelling, interactive storytelling etc.) that could help us achieve better experiences for the player with less cost/effort. These AI roles can be used in conjunction with NPC AI or just on their own.

More details on this view: see reference on the slide.

AI and Games: A bit of history – **Industry**

[see Section 1.2.2 for more details]

PacMan (Namco, 1980)

- Implicit cooperation
- Differing Ghost "Personalities"
  - Aggressive **Blinky**
  - Cunning **Pinky**
  - Wired **Inky**
  - Scared **Clyde**
- Attack – don't care waves
- No randomness

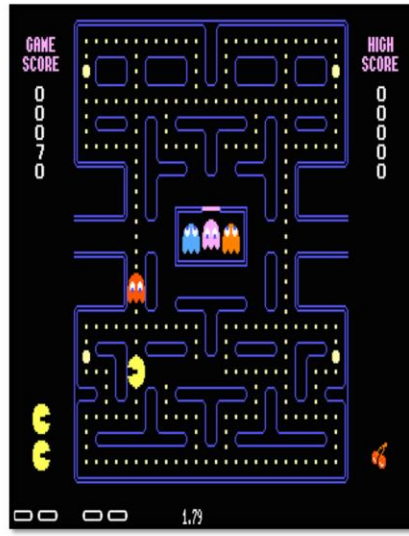

Ms Pac-Man was recently solved though – almost 1 million points through hierarchical RL

Rogue (AI Design, 1980)

- Level generation
- Constructive approach
- Inspired rogue-like game genre

Elite (Acornsoft, 1984)

- Galaxy generation
- Compression

SimCity (Maxis, 1989)

- Cellular automata
- Influence Maps

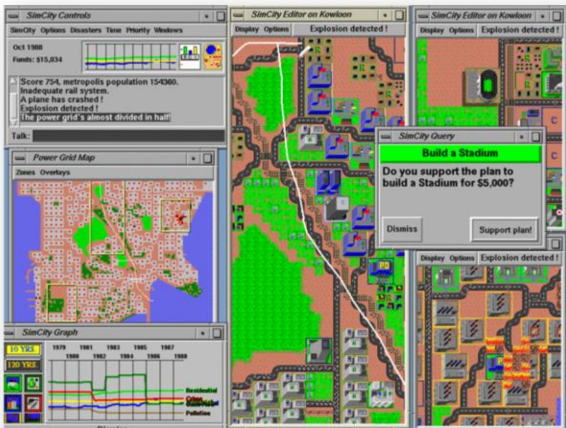

Civilization (MicroProse, 1991)

- Map/Level generation

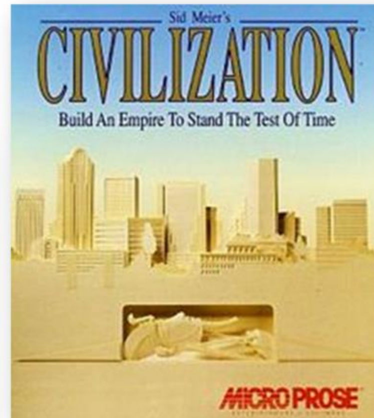

Thief (EIDOS, 1998)

- Stealth FPS, FP-Sneaker
- Advanced sensory system
- Guards perceive and respond to the environment
  - Noise, lighting, movement, and shadows.

Half-Life (Valve, 1998)

- FPS + puzzle solving
- Opponent Tactics
  - Coverage, dodging, full-level navigation
- Integration of AI into the story
- Valve hired Quake AI Mod developers for HL 2

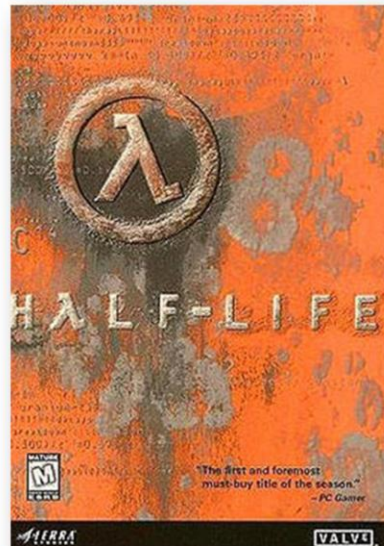

Sims (EA, 2000)

- Smart (interactive) terrains
- Smart (interactive) objects

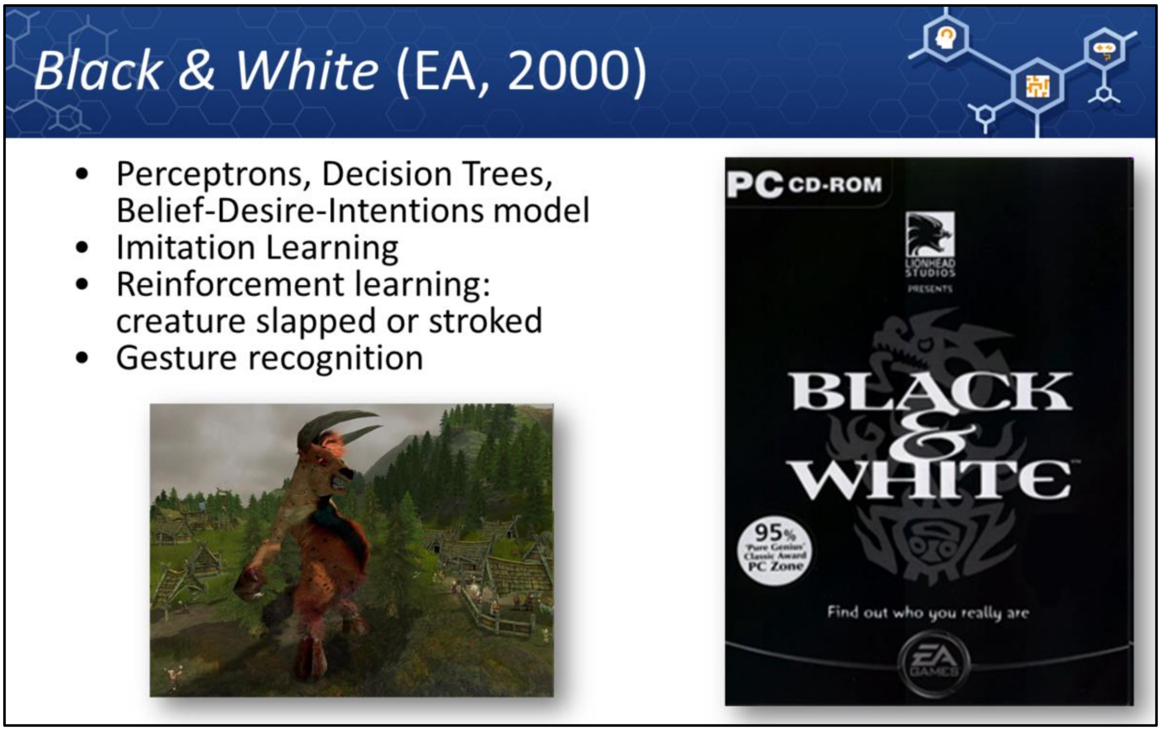

- A creature that learns through positive rewards and penalties in a reinforcement learning fashion.
- The creature employs the belief-desire-intention model for its decision making process during the game.
- The desires of the creature about particular goals are modeled via simple perceptrons.
- For each desire, the creature selects the belief that it has formed the best opinion about; opinions, in turn, are represented by decision trees.

Halo 2 (MS Game Studios, 2004)

- The revolution of **Behavior Trees**
- More human-like bots

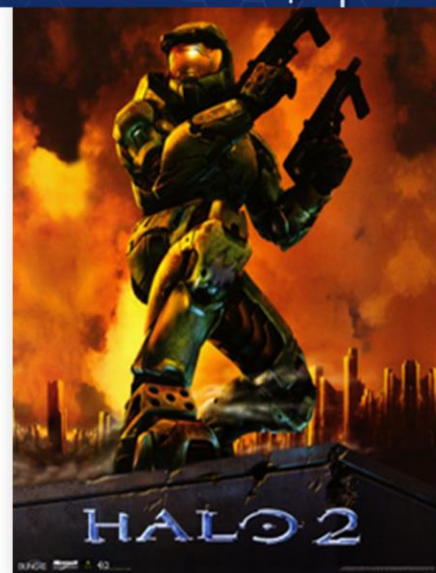

F.E.A.R. (Sierra, 2005)

- Best AI opponent – bot (at the time)
- Bot squads, teammate behaviors
- Cover when hit from behind
- Intelligent tactics
- Move through walls and under objects
- Planning system based on STRIPS
- "Environments to showcase the AI"

Forza Motorsport (MS game studios, 2005)

- Drivatar system
- Imitation Learning
- MS Research – Cambridge

Façade (Mateas & Stern, 2005)

- Interactive Storytelling
- Natural Language Processing
- AI faults: absorbed by game (character) design

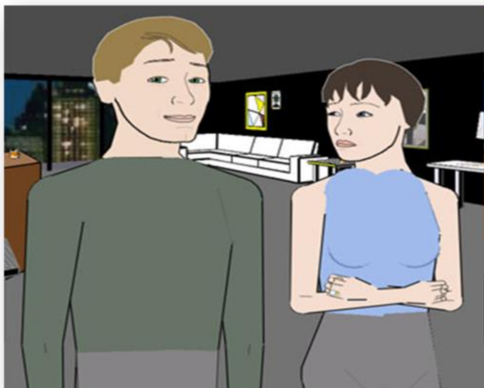

Left 4 Dead (Valve, 2008)

- **AI director**
  - Intelligent enemy placement
  - Intelligent item placement
  - Player emotional cues (visual effects, dynamic music)

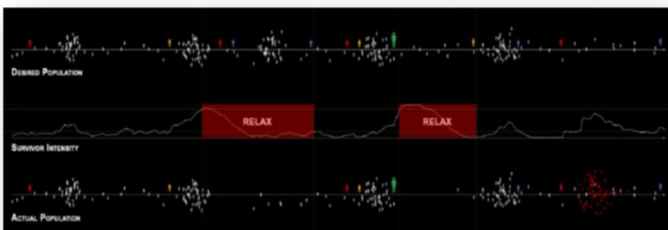

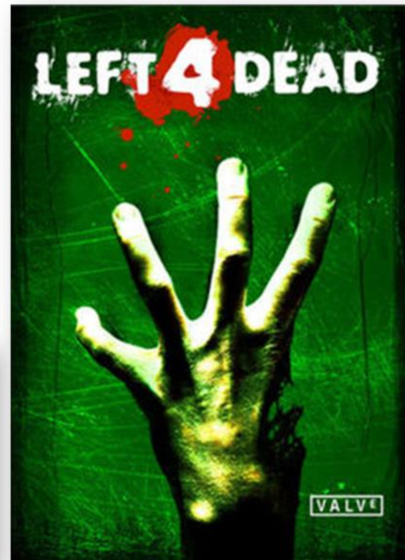

Spelunky (Mossmouth, 2008)

- Level generation
- Ultimate re-playability
- Guarantees playability

Silent Hill (Konami, 2010)

- Survival horror
- Personality-based adaptation

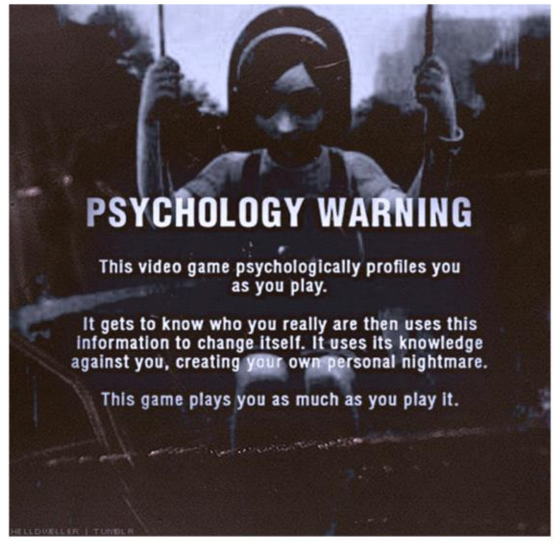

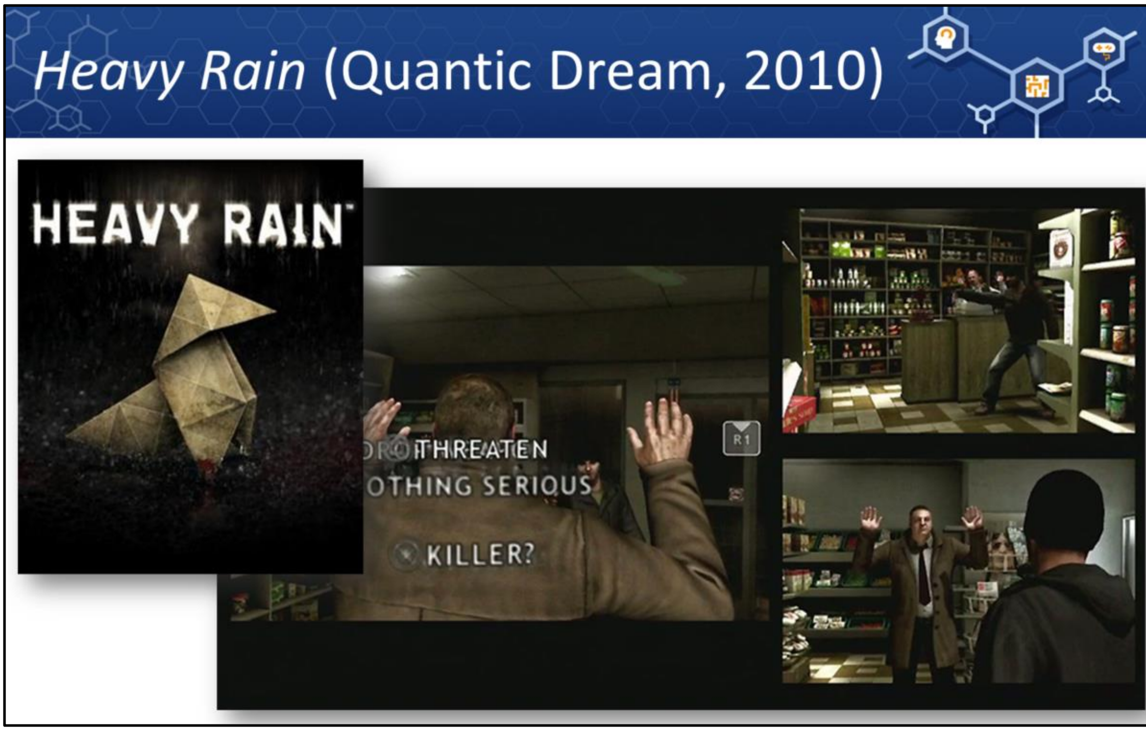

Affect-based cinematographic representation of multiple cameras

Various Kinect Games (2010 – ...)

- E.g. Kate & Milo (demo), Kinectimals, etc..
- Affective AI
- Posture (fully body) Recognition
- Speech recognition

Bioshock Infinite (2k Games, 2013)

The companion character, Elizabeth, in BioShock Infinite (2k Games, 2013);

Elizabeth, a crucial element of the game, was designed as a character which could not only be a useful AI companion to the player but a real partner with a significant emotional bond as well

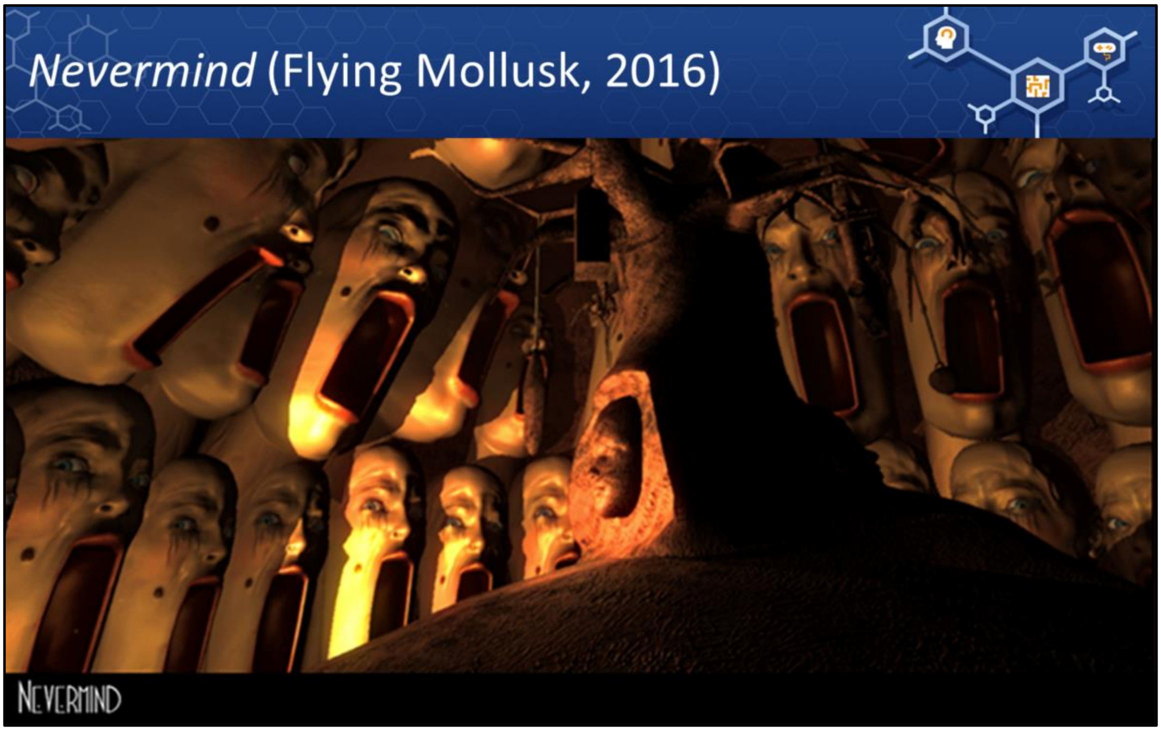

Survival horror - affect-based game adaptation via a multitude of physiological sensors.

No Man's Sky (Hello Games, 2016)

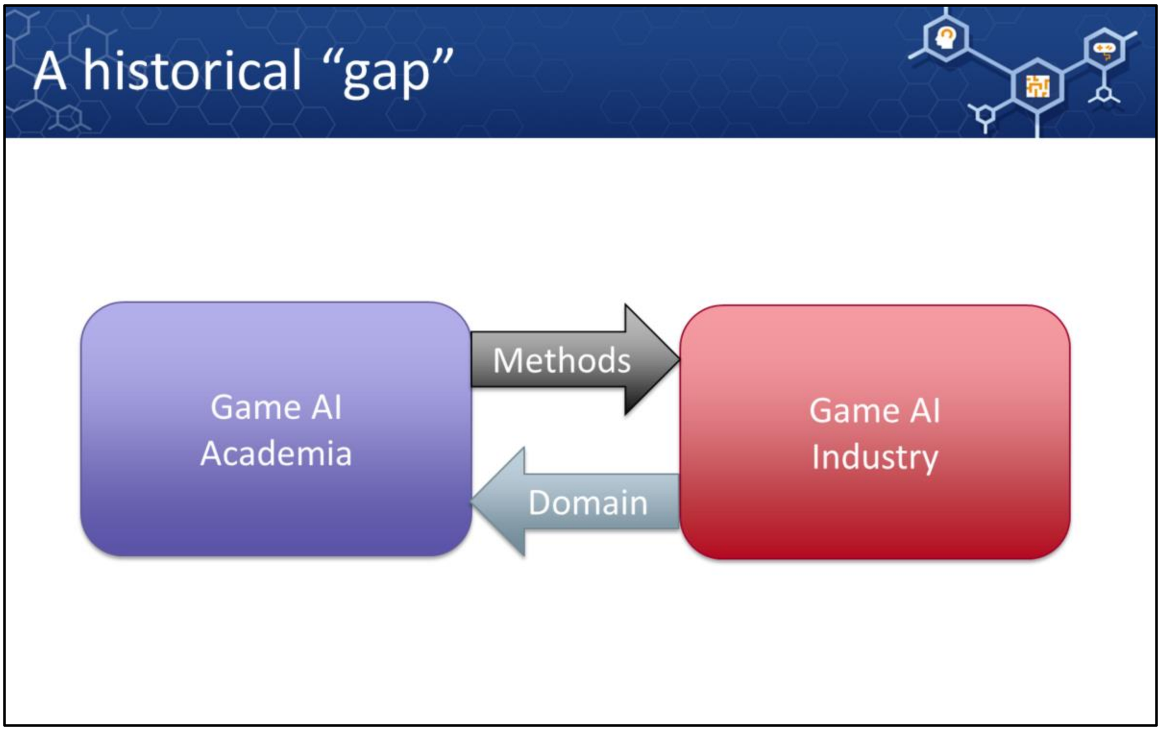

[See Section 1.2.3 for more details]

Some traditional "complaints" – myths to a large degree

- Academia: Industry does not use our tools
- Industry: Academics lack domain knowledge and ignore our problems
- Game AI researchers do not collaborate with industry

- Not much of a gap after all
- The gap is there in problems that academics do not wish to solve and industry does not car about. But that is a healthy one – more progress need to be made

- Industry: NPC AI is more or less "solved"(e.g. look at Left 4 Dead, F.E.A.R. and Skyrim; Satisfying NPC behaviors)
- AI tracks in Game Developers Conference do not necessarily focus on NPC AI
- Academic Panels/Special Sessions/Tutorials (FDG, CIG, AIIDE), Special Issues (IEEE TCIAIG/ToG) focus on non-traditional uses of AI
- More effective and active industry/academia communication/collaboration over the years
- Academic support on multidisciplinary nature of game AI
- Pragmatic and holistic view of game AI

Games helps us design **better** AI…

…and AI can help us make **better** games ...

…but not necessarily through **game agents**!

How then?

G. N. Yannakakis, “**Game AI Revisited**,” in *Proceedings of the 9th ACM Computing Frontiers Conference*, pp. 285-292, 2012. (**invited paper**)

Overall summary: Game AI is not all about NPC AI

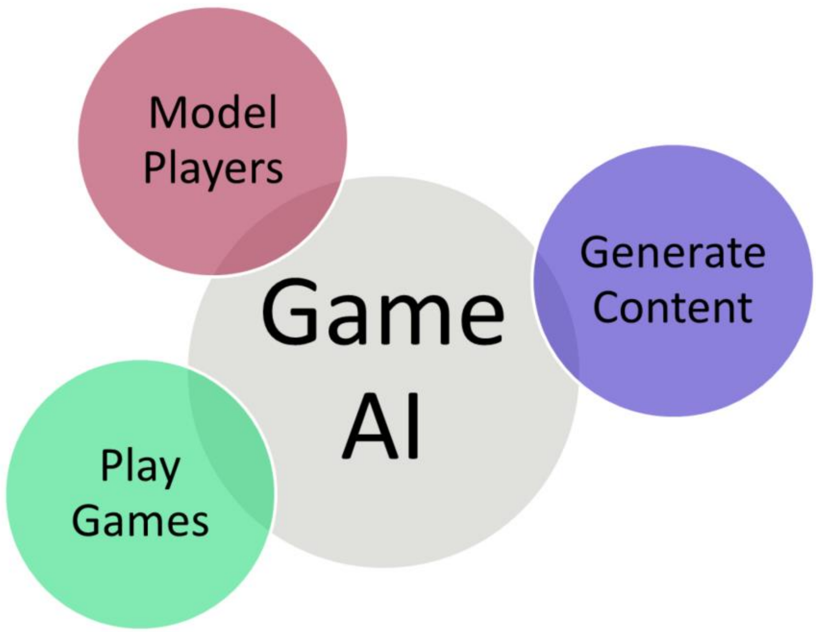

G. N. Yannakakis and J. Togelius, “**Artificial Intelligence and Games**,” Springer Nature, 2018

The tree main roles that AI takes *in* and *for* Games

- Play games [Chapter 3]
- Generate Content [Chapter 4]
- Model Players [Chapter 5]

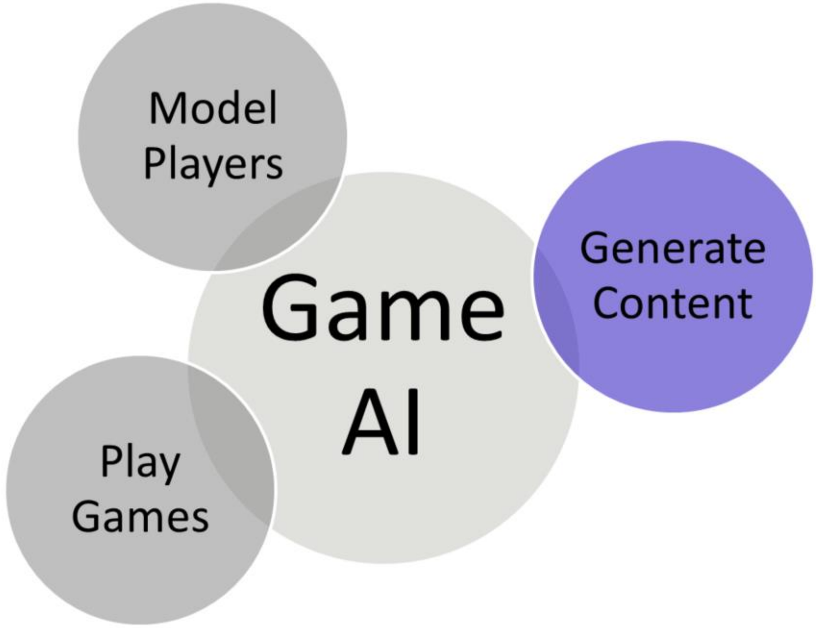

G. N. Yannakakis and J. Togelius, “**Artificial Intelligence and Games**,” Springer Nature, 2018

A brief introduction to the “Generate content” role

Procedural Content Generation

What is PCG? "The creation / production of new game content (semi) automatically via algorithmic means "

Why is PCG important?

- Lowers development costs
- Enables adaptability
- Increases replayability
- Design beyond human creativity (?)

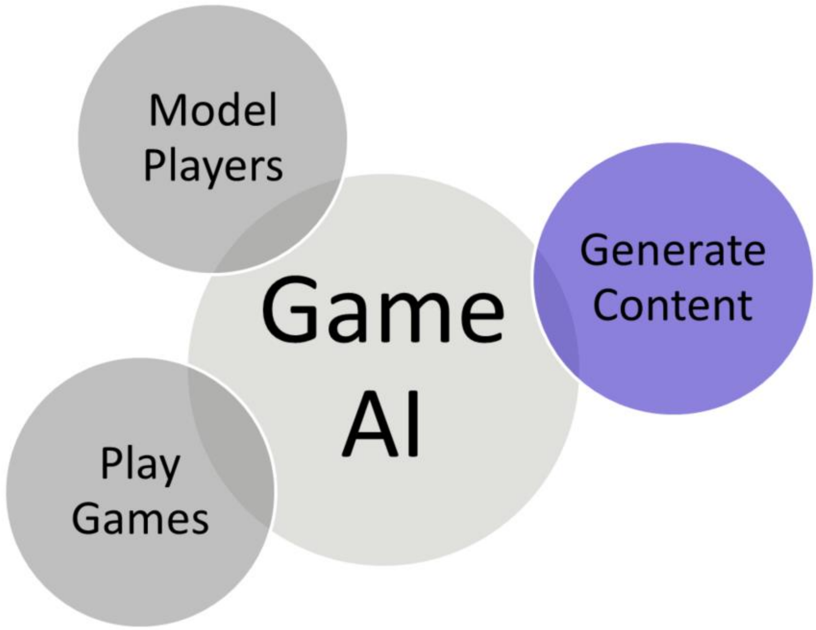

G. N. Yannakakis and J. Togelius, "**Artificial Intelligence and Games**," Springer Nature, 2018

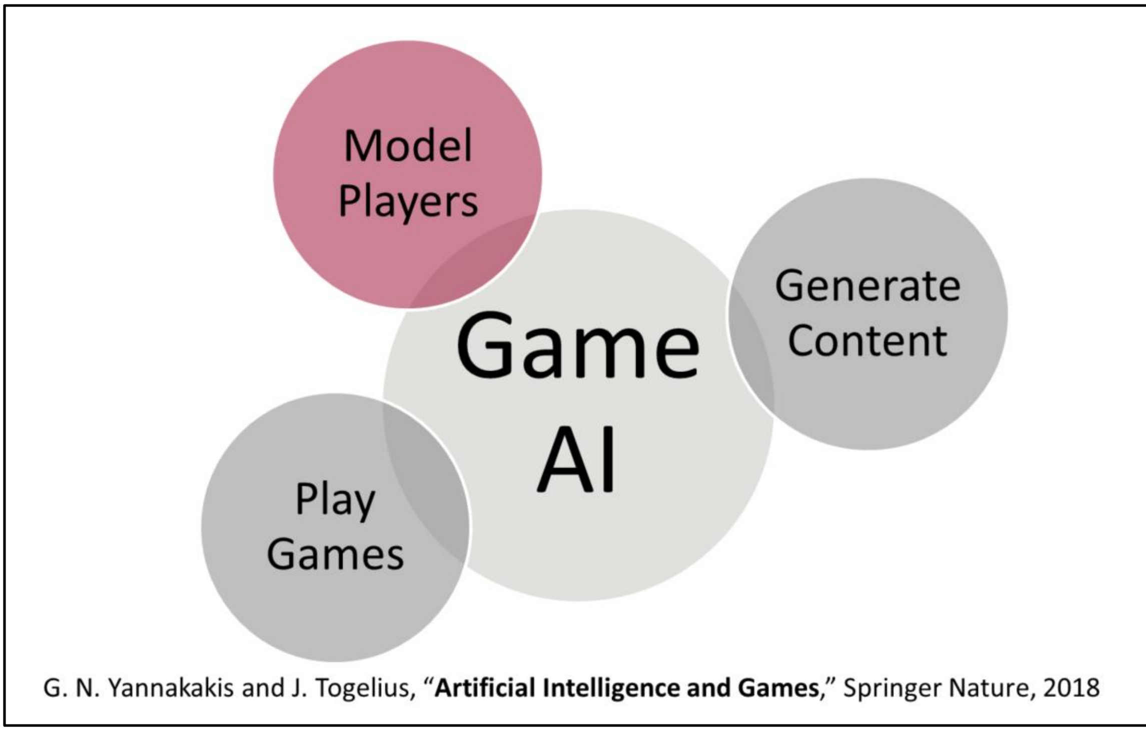

Now onto a brief introduction to the "model players" role

“Games cannot be dissociated from **emotion** and **learning**…”

G. N. Yannakakis and A. Paiva “**Emotion in Games**,” in *Oxford Handbook of Affective Computing*, 2014

Why is there an association between games, emotion and learning?

- We play games throughout our lives – we learn
- We experience positive negative (and even) emotions voluntarily! Impressive!
- We can control our experiences via adaptation (not like tv, filmes etc..)

Why Player Modeling is Important?

- There is no game without a player…
- The perfect suit is tailored to you! How about the perfect game?
- We (players) are very different
- We (players) are many more than before
- **Holy grail** of game development

“Games: the **final frontier** for AI?”

“AI: the **next step** for Games!”

Summary of the Introduction

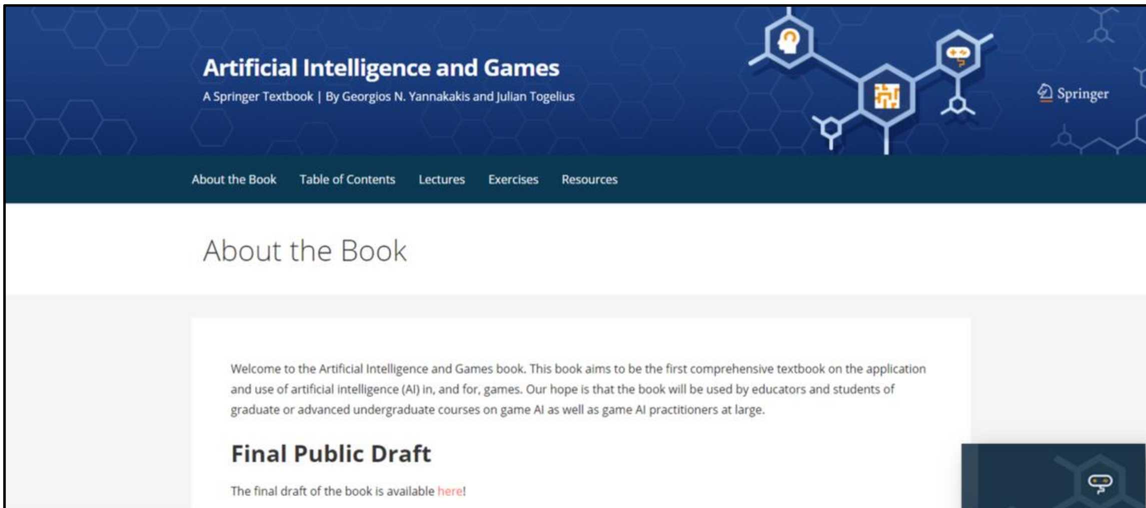

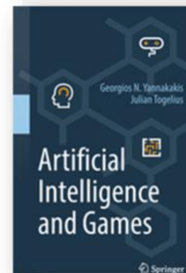

Readings: Chapter 1

gameaibook.org