SEED // SEARCH FOR EXTRAORDINARY EXPERIENCES DIVISION

ML in Game Production

Linus Gisslén – ML Technical Director
SEED – EA

Electronic Arts

August 2023 - Conference on Games

Purpose with this talk

- To show **interesting challenges and solutions** using ML in games and game production.
- To show why **ML is, and will continue to be important** for game production in the future.
- Present some **opportunities** for research in this domain and to foster even more **collaboration** between academia and industry.
- **Excite** you about the possibilities to work with this!

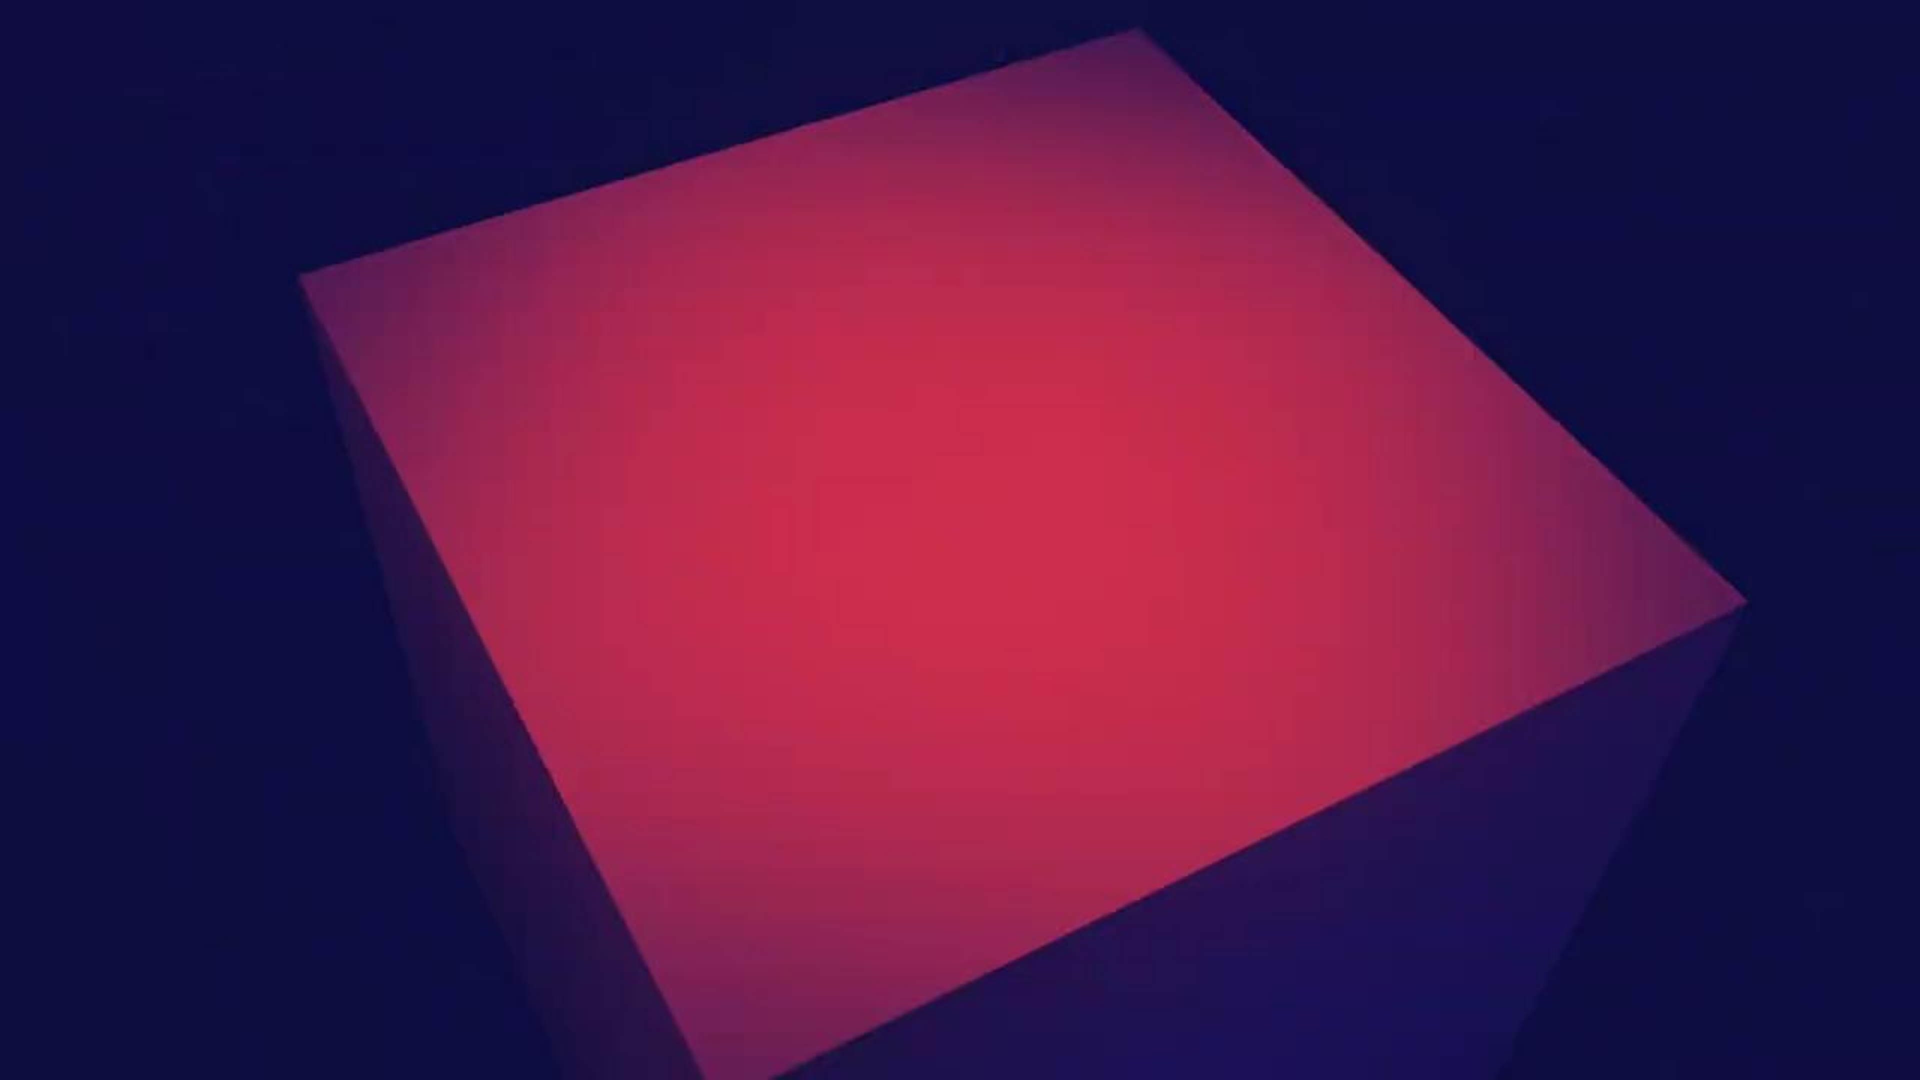

Electronic Arts

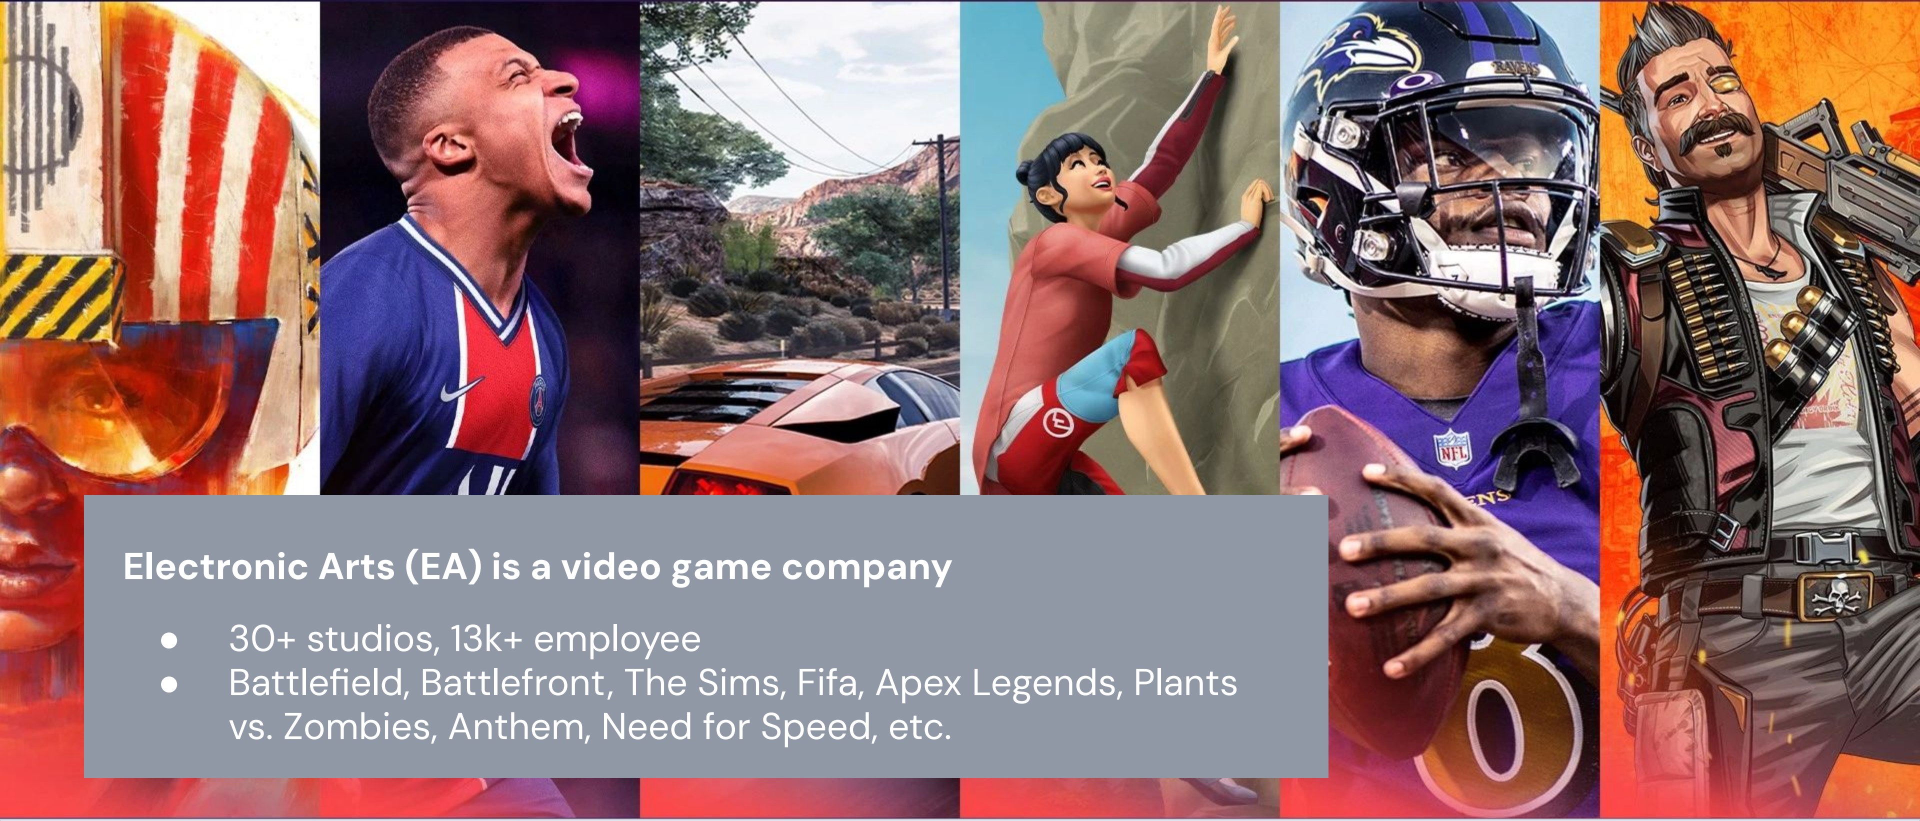

Electronic Arts (EA) is a video game company

- 30+ studios, 13k+ employee
- Battlefield, Battlefront, The Sims, Fifa, Apex Legends, Plants vs. Zombies, Anthem, Need for Speed, etc.


SEED

SEED

Search for Extraordinary Experiences Division

SEED is an advanced R&D department at EA. Our mission is ***to explore, build, and help define the future of interactive entertainment***

Locations: Stockholm // San Francisco // LA // Vancouver // Montreal // Remote

We exist to:

- Explore risky topics
- Research new technologies for innovative concepts to emerge
- Empower EA to create even greater experiences for our players

EA ad from 1983

About SEED: Our different research vectors

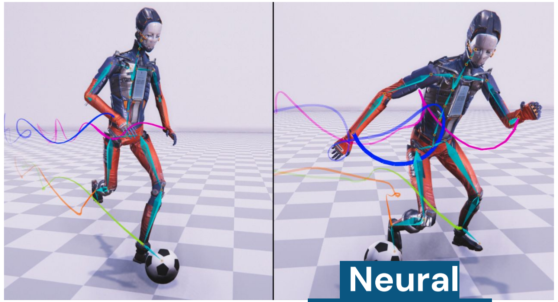

SEED partnerships

Current academic partnerships:

- BAIR: Pieter Abbeel, Trevor Darrell, Angjoo Kanazawa
- KTH: TMH, Robotics, Digital Future, etc.
- University of Houston
- University of Florence: Andy Bagdanov
- WASP WARA Media & Language

Open for suggestions! :)

Bridge between academia and industry.

Industrial partnerships

ML for game production

Motivation for ML in game production

Today's games are often huge and requires years of development

Case study: Red Dead Redemption 2

- Cost: ~$500M
- Time: ~8 years of development
- Staff: ~2000 persons

Trend in games

Trend in AAA games:

- Open world
- Higher quality assets, more detail
- Deeper interactions

Trend continues for industry and EA. Automation and tools to create new generation of games is crucial.

ML is ideal for this and in this talk I will give some example of research in this direction.

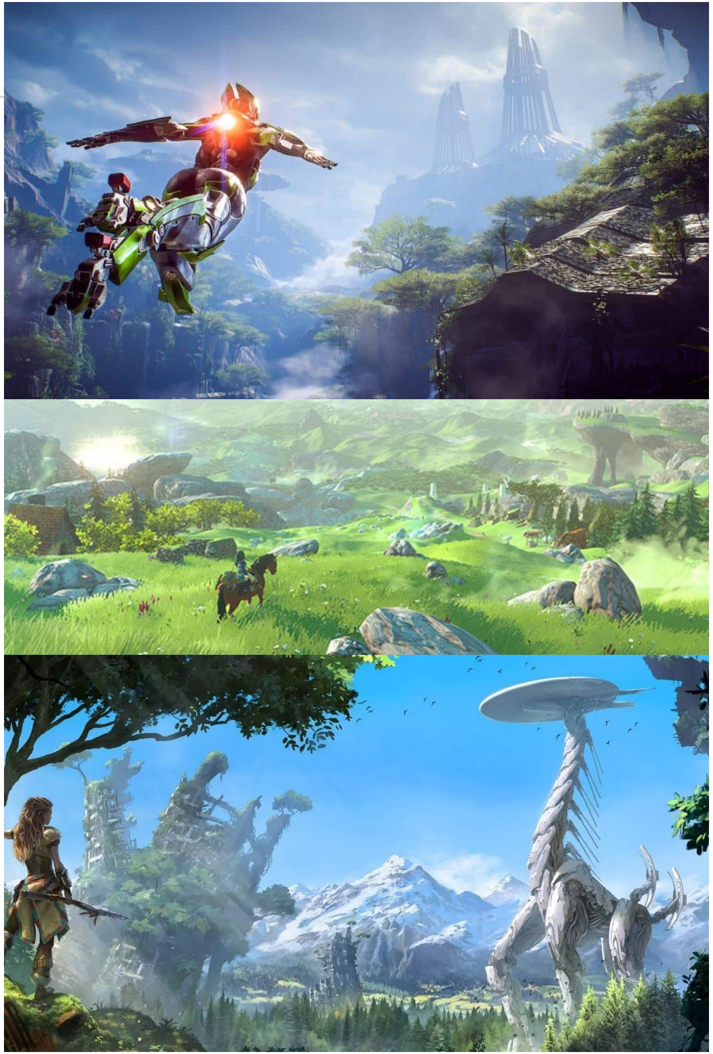

Data Driven Content Generation

Automated facial animations

Motivation

- Creating facial animations is very tedious, takes days to do minutes of animations
- Existing automated (non-ML) methods are not very good
- E.g. Real-time animation on player speech would unlock new experiences

Voice2Face

Generating Facial Expressions from Voice

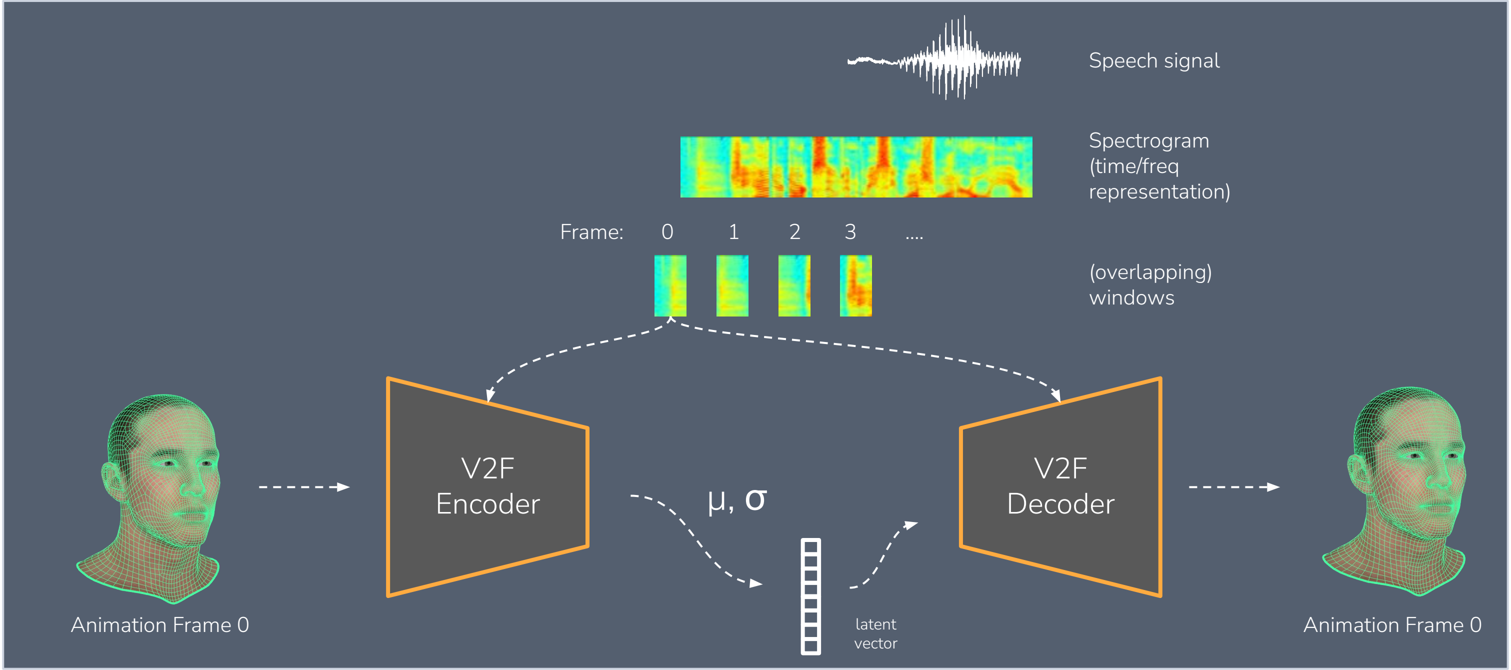

Voice2Face

Results: English

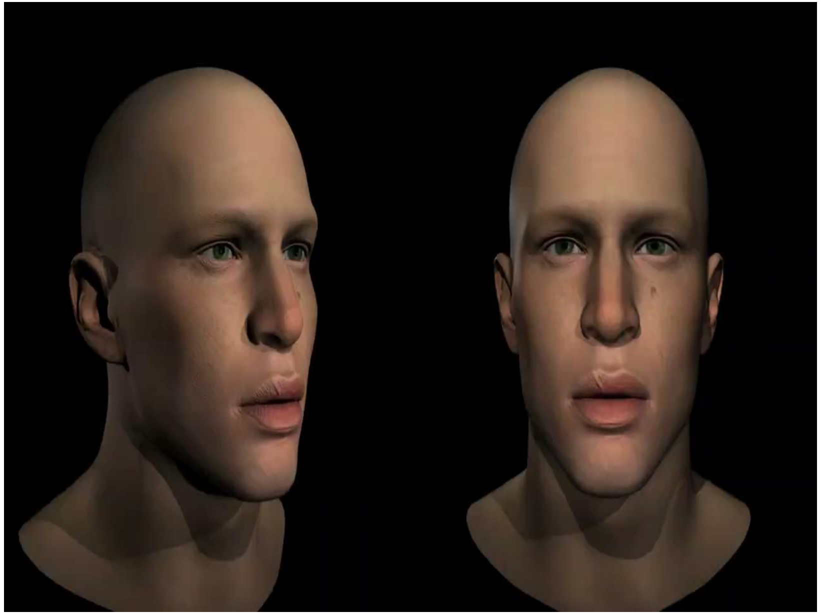

Non-English

Voice2Face: Audio-driven Facial and Tongue Rig Animations with cVAEs

Aylagas MV, Leon HA, Teye M, Tollmar K Computer Graphics Forum 2022

Expressive Voice2Face

Comparison, different emotions

Cloth Simulation

Swish: Neural Cloth Simulation on Madden 21-24

Adding wrinkles in Jerseys with ML

- Training data:
  - Plausible poses extracted from game
  - Cloth generated in Marvelous Designer was paired with the pose. Post-processing with Maya to optimize for speed.
- Train a simple neural network pose -> mesh
- Inference time; 140 μs -> Cheaper than standard real-time cloth deformation

Visual target

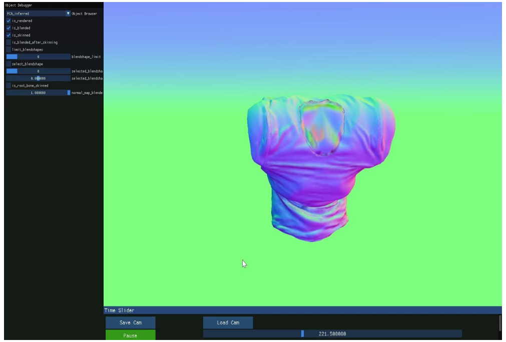

Training data sample

Swish: Neural Cloth Simulation on Madden 21-24

Adding wrinkles in Jerseys with ML

Swish: Neural Network Cloth Simulation on Madden NFL 21

Chris Lewin
chlewin@ea.com
Electronic Arts / SEED
UK

James Power
jpower@ea.com
Electronic Arts / Tiburon
USA

James Cobb
jcobb@ea.com
Electronic Arts / Tiburon
USA

Figure 1: Player jerseys in Madden NFL 21 simulated using our system.

Swish: Neural Network Cloth Simulation on Madden NFL 21.
C. Lewin, J. Power, J. Cobb. ACM SIGGRAPH 2021.

Results in-game (Madden 21)

Swish: Neural Cloth Simulation on Madden 21-24

Adding wrinkles in Jerseys with ML

Potential future research:

- More loosely fit garment: dresses, capes, scarves, etc.
- Other cloth objects such as flags and tents
- Other types such as trees, grass, waves, etc.


Results in-game (Madden 21)

Swish: Neural Network Cloth Simulation on Madden NFL 21.
C. Lewin, J. Power, J. Cobb. ACM SIGGRAPH 2021.

Self Learning Agents

Self-Learning Agents

Data driven behaviour generation

Benefits over traditional methods (scripting):

- No scripting/programming needed: more accessible to everyone
- Re-train instead of re-script: Less manual work
- Automated (can run overnight): potentially faster
- Can solve problems that scripting can not: adds functionality

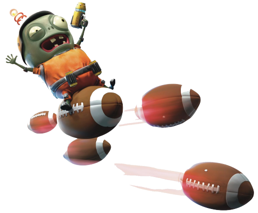

Self-Learning Agents for Game AI

Some lessons learned from adding RL to Battlefield 2018:

- Artists needs explicit control over behaviours/visuals
- Inference (run-time) took too much resources
- Player facing ML is difficult, no room for failures
- It did not fit in into the existing pipeline of game creation
- Expertise on the “receiving” end did not exist, i.e. hard to retrain when needed

**RL agents playing Battlefield 1*

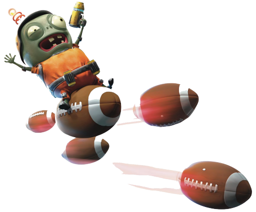

**Imitation Learning with Concurrent Actions in 3D games, J Harmer, et. al CoG 2018*

Self-Learning Agents for Game Testing

✅ No artist needs to see the failures
✅ No run time needed, at least not critical
✅ Not player facing
❌ Needs to fit into existing pipeline
❌ Expertise on the “receiving” end needed

3/5 is not too bad odds!

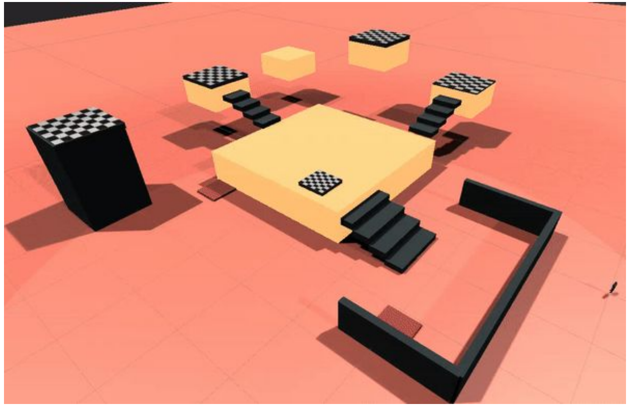

Test-bed for game testing algorithms

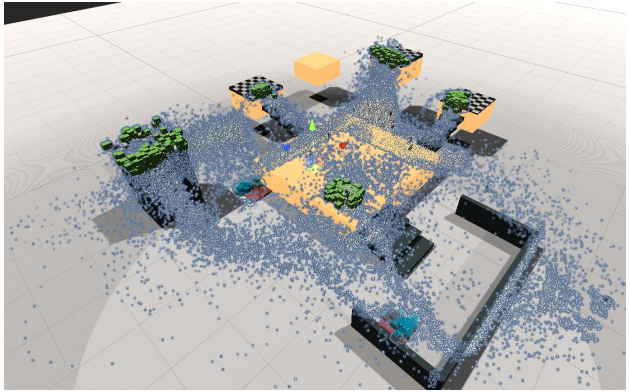

First use-case: Automated Game Testing

"In Battlefield V testing all maps and modes for 1 hour requires 2304h of testing. ***288 people*** *to test that every day. If we add more maps and modes this number will be larger." AI for Testing: The Development of Bots that Play 'Battlefield V'* Jonas Gillberg - GDC 2019

Problem statement:

- Automatic testing is difficult, requires scripting of bots
- Large game and open world games with procedural content doesn't scale well with current solutions
- Testing of user generated content requires new solutions, something that can adapt

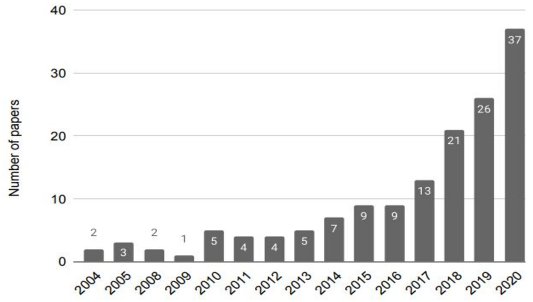

*Trend in research: exponential growth in published papers on "Automated game testing"**

**Towards Automated Video Game Testing: Still a Long Way to Go*
C. Politowski, YG Guéhéneuc, F Petrillo, ICSE Workshop on Games 2022

Automated Game Testing with Reinforcement Learning (RL)

Goal for RL is to maximize reward -> exploitation

Finding exploits without being "told" (see right)

Algorithm learns, therefore:

- No scripting required
- Retrain instead of rewriting scripts
- Unpredictable and teachable: more "human like" control
- Explorative: by figuring out how to play it covers more game states

Courtesy: Open AI blog *Faulty reward functions in the wild*


Trained for 24 hours

Automated Game Testing

High game state coverage -> more bugs found -> better testing. A RL case study.

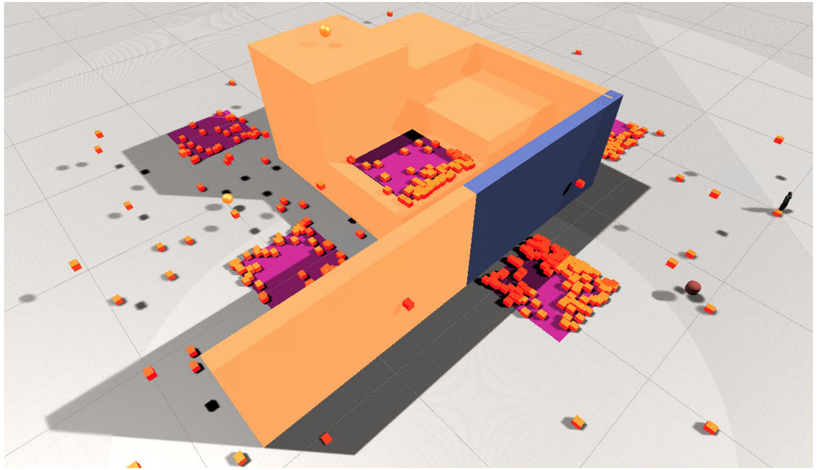

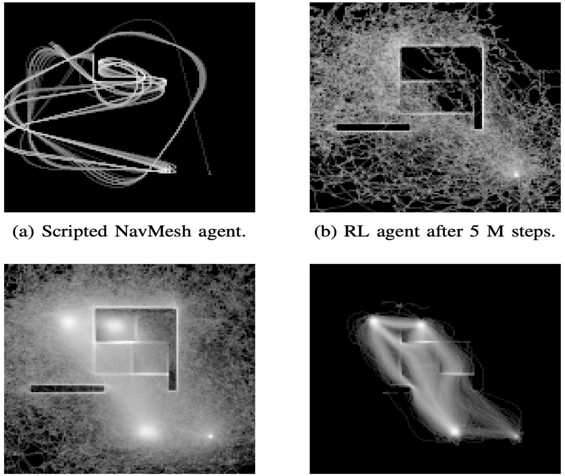

(a) Scripted NavMesh agent. (b) RL agent after 5 M steps.

(c) RL agent after 30 M steps. (d) RL agent fully trained.

Augmenting automated game testing with deep reinforcement learning
J Bergdahl, C Gordillo, K Tollmar, L Gisslén. CoG 2020

Problem: RL does not always generalize well

Inspiration:

1. Idea: That training on procedurally generated content improves generalization in RL agents*
2. Idea: That RL can be used for PCG**
3. Idea: Posing increasingly difficult (progressive PCG) problems increase learning capacity***

Proposed solution: Generator RL and Solver RL pair. With the feedback from the Solver the Generator can learn to make difficult but not impossible maps.

* *Increasing Generality in ML through PCG*. Risi S & Togelius J Nature Machine Intelligence 2020
** *Pcgrl: Procedural content generation via reinforcement learning* Khalifa A, Bontrager P, Earle S, Togelius J. AIIDE 2020
*** *Illuminating generalization in deep reinforcement learning through procedural level generation* Justesen, Niels, et al. ArXiv 2018

ARLPCG: Adversarial RL for Procedural Content Generation

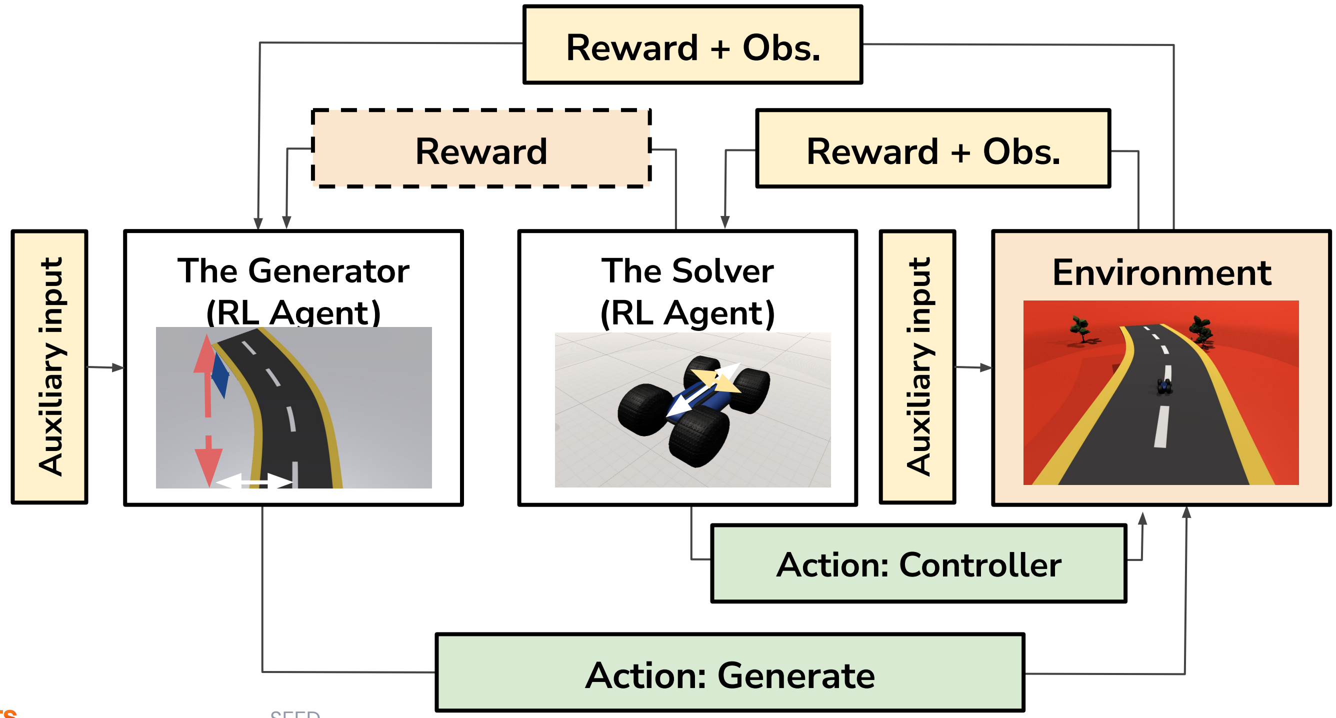


y
x
Persp
-1,0
Adversarial RL Agents
Generator creates the environment / Solver explores the environment

Automatic Gameplay Testing with Curiosity-Conditioned Proximal Trajectories

Combining Imitation Learning and Curiosity for guided exploration.


+

=


Automatic Gameplay Testing with Curiosity-Conditioned Proximal Trajectories

Combining Imitation Learning and Curiosity for guided exploration.

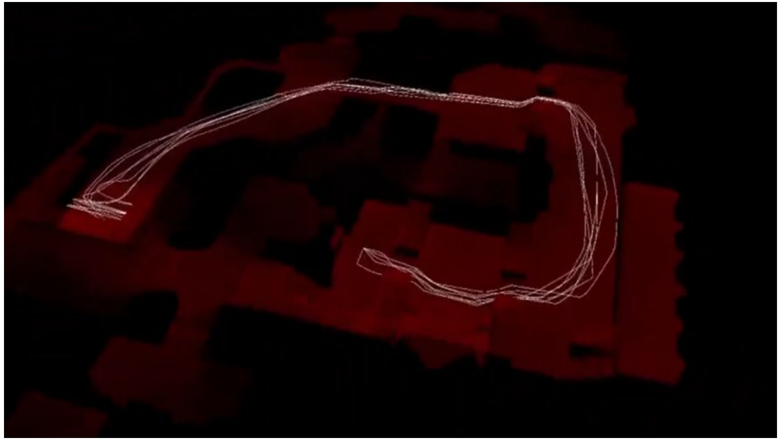

Human demonstration in VizDoom

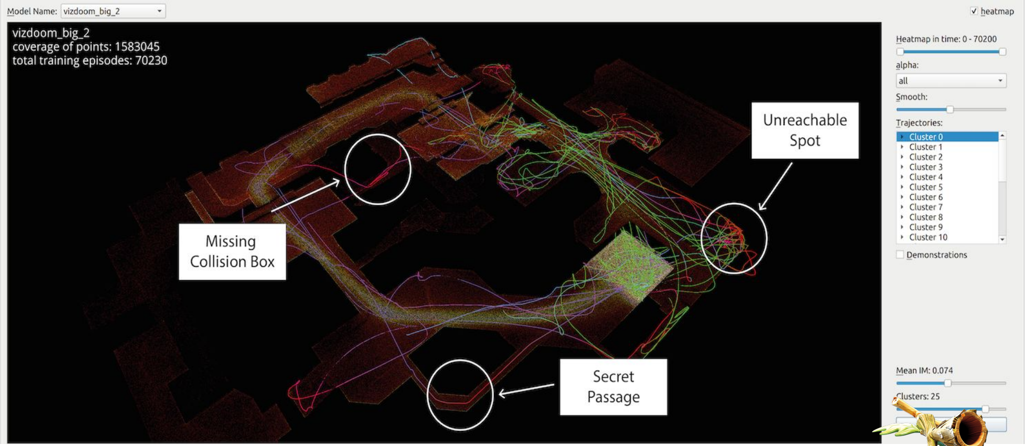

Algorithm exploration includes a unknown passage

Electronic Arts

Automatic Gameplay Testing with Curiosity-Conditioned Proximal Trajectories

Automated Gameplay Testing and Validation with Curiosity-Conditioned Proximal Trajectories
A Sestini, L Gisslén, J Bergdahl, K Tollmar, AD Bagdanov, Transactions on Games 2022

Technical Challenges of Deploying Reinforcement Learning Agents for Game Testing in AAA Games

- Time: 13:50, Wednesday August 23rd. Location: ISEC 102

Technical Challenges of Deploying Reinforcement Learning Agents for Game Testing in AAA Games
J Gillberg, J Bergdahl, A Sestini, A Eakins, L Gisslén. Conference on Games (CoG) 2023

Self-learning Agents for Game Design

The Problem

- Game testing is today often a long process with long lead times. Testing happens often in another phase than design phase.
- Test results comes too late for it to be efficient use of data, and sometimes too late for corrections to be added in production.

Proposed solution

- Let expert (game designers) demonstrate how level should be played.
- ML agents learn from that data, and plays the game accordingly.
- Move the testing “upstream” i.e. in the hands of the creators/designers.

Self–Learning Agents for Game Design

Requirements for this use-case

- End-user (game designers, etc.) does not know ML (so reward shaping is difficult)
- No, or little, exploration and exploitation is desirable
- Speed: Has to be fast enough to be used in production
- Generalization: Small changes should not lead to re-training
- Controllability: Human Personas is important in game design. This should be incorporated.

Imitation Learning for Game Design

Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning

- 11:00, Thursday August 24th. Location: ISEC 140

Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning

Alessandro Sestini[1], Joakim Bergdahl[1], Konrad Tollmar[1], Andrew D. Bagdanov[2], Linus Gisslén[1]
[1]*SEED - Electronic Arts (EA)*, [2]*Università degli Studi di Firenze*
{asestini, jbergdahl, ktollmar, lgisslen}@ea.com
andrew.bagdanov@unifi.it

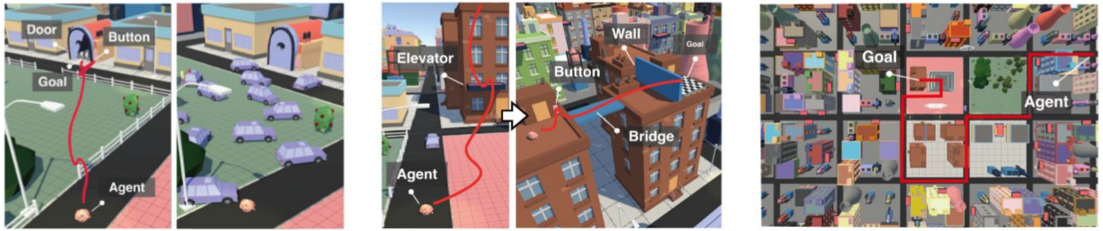

Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning
A Sestini, J Bergdahl, A Bagdanov, K Tollmar, L Gisslén. Conference on Games (CoG) 2023

Generating Personas for Games with Multimodal Adversarial Imitation Learning

- 09:40 Thursday August 24th (Best of CoG). Location: ISEC 102

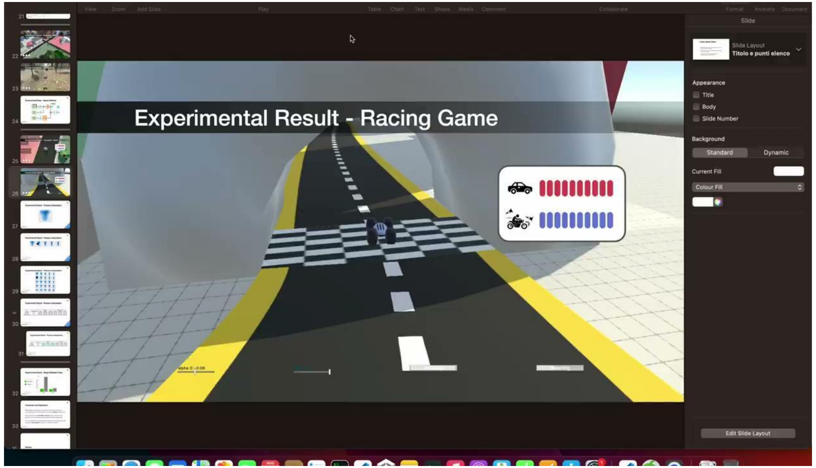

Generating Personas for Games with Multimodal Adversarial Imitation Learning
W Ahlberg, A Sestini, K Tollmar, L Gisslén Conference on Games (CoG) 2023

Potential use cases/research for ML in games

Future of ML in games and game production

Caveat

My prediction (in 2011, when I started in AI):

2025	2050
Unloading dishwasher Cleaning Self-driving cars	Poetry Art Music

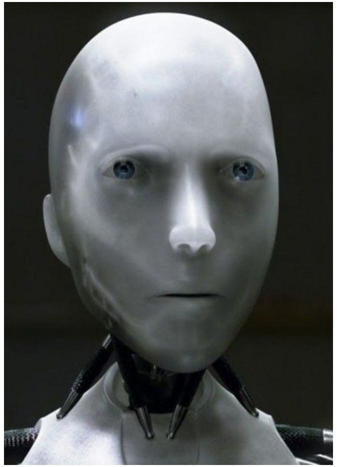

Caveat

The NeverEnding Game: How AI Will Create a New Category of Games

by Jonathan Lai

Reality (now):

2025 → 2050

Poetry
Art
Music

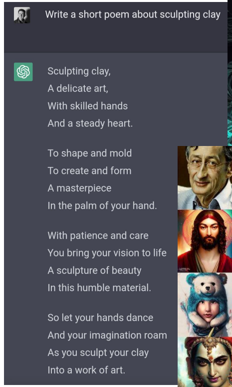

Unloading dishwasher
Cleaning
Self-driving cars

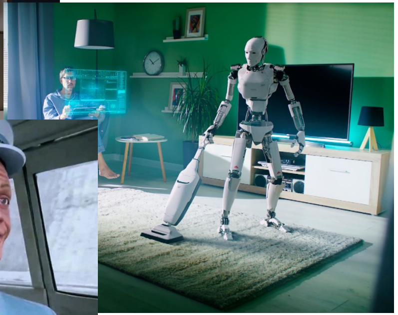

Future of ML in games and game production

Dialogues with NPCs

GPT4 for NPCs demo from May 2023
Courtesy: NVIDIA

Future of ML in games and game production

Possible use-cases for LLMs:

- NPC dialogue (see video)
- Player coach
- Brainstorm tool
- Side-quest generator
- Game commentator
- Etc. etc.

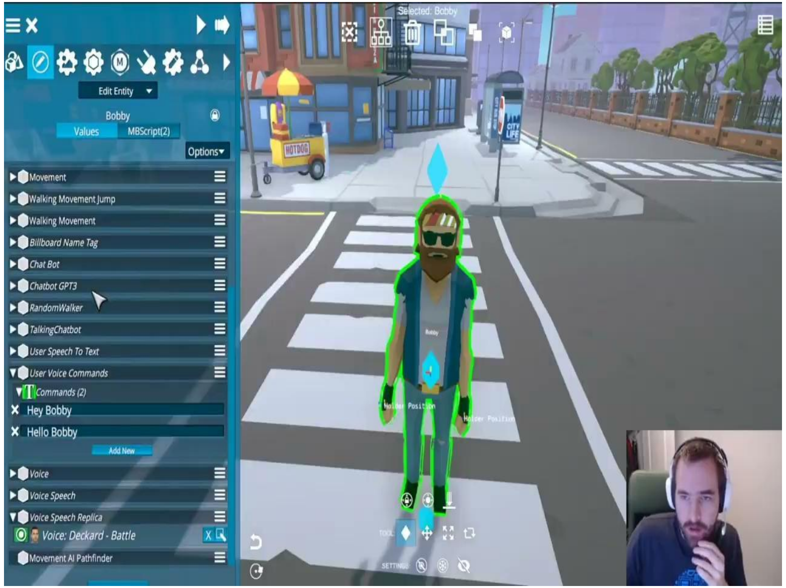

Use case: NPC dialogue (pauses are cut).
GPT3 demo from Feb 2021.
Courtesy: Lee Vermeulen

Future of ML in games and game production

- The industry is moving from manual and scripted to data driven methods
- Data driven content generation + data driven behaviour.
- As ML tools becomes more available -> domain knowledge will be more important

Courtesy: Midjourney AI Art

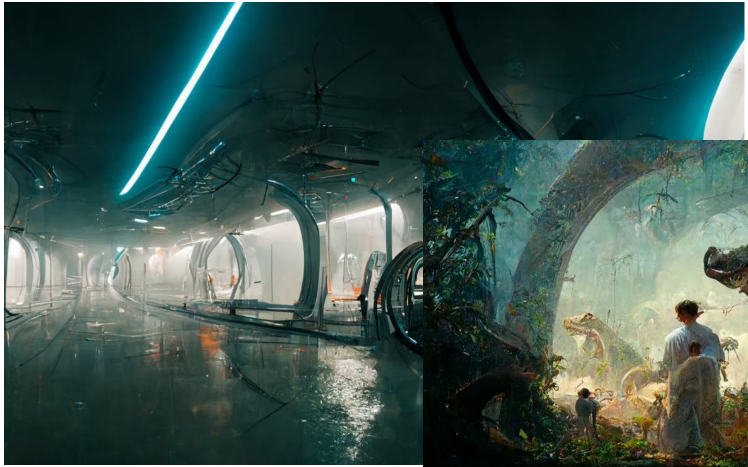

Future of ML in games and game production

Challenges for data driven solutions

- **Curation:** how do we filter? Guard-rails for player facing ML.
- **Speed:** Generation might be expensive and slow, especially in real-time applications.
- **Control:** How can we control the output and make sure it fits with the narrative
- **Data:** Acquisition or generation of data is a non-trivial process and highly affects the quality of the models

SPECULATIVE: Game creation will change significantly in 5-10 years

- As algorithm becomes more "intelligent" game devs and designers role will become more a of a movie director rather than to explicit instruct/create
- Devs will be able to *instruct by natural language and show by examples*

Current state

- Programming of behaviours
- Largely manually creation of assets

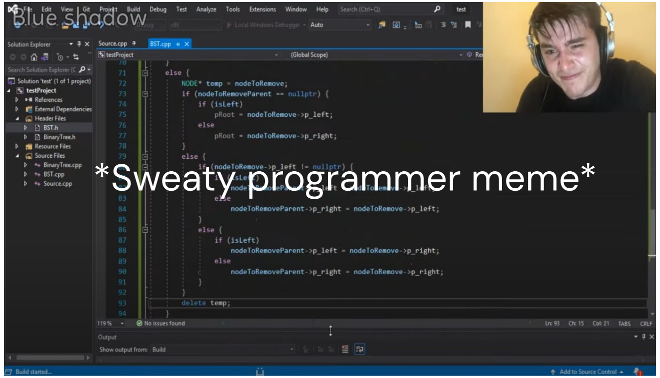

"If(playerDistance < attackDistance) Attack(player);"

Future

- Demonstrating behaviours
- Showing examples to generate new assets
- Instructing through natural language

"Attack the player when they come too close"

Summary

Summary

We believe that ML and AI has the potential to radically **change the way we create games**

So far the industry have **only explored a fraction** of what's possible using ML/AI

Change is coming to game production, you better **be prepared** :)

Thank you for listening!

seed.ea.com

Contact:
Linus Gisslén (lgisslen@ea.com)

References

1. *Generating Personas for Games with Multimodal Adversarial Imitation Learning W Ahlberg, A Sestini, K Tollmar, L Gisslén Conference on Games (CoG) 2023*
2. *Technical Challenges of Deploying Reinforcement Learning Agents for Game Testing in AAA Games J Gillberg, J Bergdahl, A Sestini, A Eakins, L Gisslén. Conference on Games (CoG) 2023*
3. *Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning A Sestini, J Bergdahl, A Bagdanov, K Tollmar, L Gisslén. Conference on Games (CoG) 2023*
4. *Neural Synthesis of Sound Effects Using Flow-Based Deep Generative Models* S Andreu, M Villanueva Aylagas AIIDE 2022
5. *Voice2Face: Audio-driven Facial and Tongue Rig Animations with cVAEs Monica Villanueva Aylagas, Hector Anadon Leon, Mattias Teye, Konrad Tollmar*
6. *Swish: Neural Network Cloth Simulation on Madden NFL 21. C. Lewin, J. Power, J. Cobb. ACM SIGGRAPH 2021.*
7. *Imitation Learning with Concurrent Actions in 3D games, J Harmer, L Gisslén, J del Val, H Holst, J Bergdahl, T Olsson, K Sjöö, M Nordin. Conference on Computational Intelligence and Games 2018*
8. *Augmenting automated game testing with deep reinforcement learning J Bergdahl, C Gordillo, K Tollmar, L Gisslén. Conference on Games CoG 2020*
9. *Improving Playtesting Coverage via Curiosity Driven Reinforcement Learning Agents C Gordillo, J Bergdahl, K Tollmar, L Gisslén. CoG 2021*
10. *Automated Gameplay Testing and Validation with Curiosity-Conditioned Proximal Trajectories A Sestini et. al 2022*
11. *Adversarial Reinforcement Learning for Procedural Content Generation L Gisslén, A Eakins, C Gordillo, J Bergdahl, K Tollmar. CoG 2021*
12. *Automatic Testing and Validation of Level of Detail Reductions Through Supervised Learning*, Tamm et. al CoG 2022
13. *Using deep convolutional neural networks to detect rendered glitches in video games*, C. Ling, K. Tollmar, L. Gisslén AIIDE 2020