

GoodAI:

- 1) Our research
- 2) Teaching robots to walk in Space Engineers

Week of Innovation, October 22

Martin Poliak



GoodAI

Our History



FOUNDED IN 2014 BY
MAREK ROSA



KEEN SOFTWARE HOUSE

Space Engineers Game
4M+ copies sold, \$73M revenue




PERSONAL INVESTMENT
\$10M USD



FOCUS ON "GENERAL AI"
AI in games motivating mid-term goals

GoodAI Our Work



GOODAI GRANTS

800k+ USD



GENERAL AI CHALLENGE

2 Rounds, 2 Collaborations



BADGER SEMINARS

2 Seminars, 1 Workshop



EU PROJECT PARTICIPATION

VeriDream, IV4XR, AI4EU



EU POLICY CONTRIBUTIONS

EC Report, UNESCO Draft Recommendation

1) Our research

Paths to AGI

a) Groups

- Cooperation in groups achieves more than what a single individual can
- Unlocking group learning and cooperation
-> **Memetic Badger**

b) LLMs


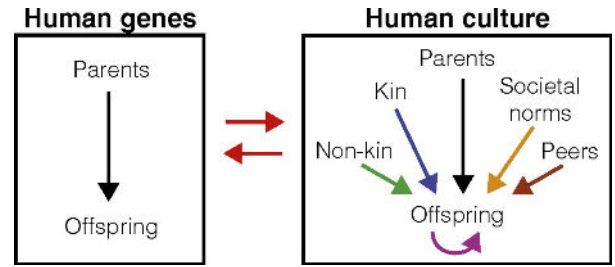

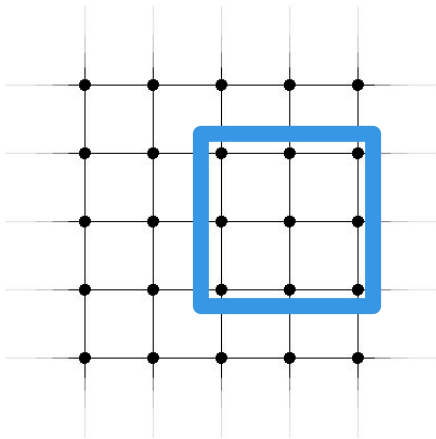
- Recent Large-Language Models (LLMs) show impressive learning capabilities, but they lack the ability to learn continually
- If we add a growing memory to them appropriately, will we be able to achieve continual learning? -> **LLM + growing memory**



Memetic Badger

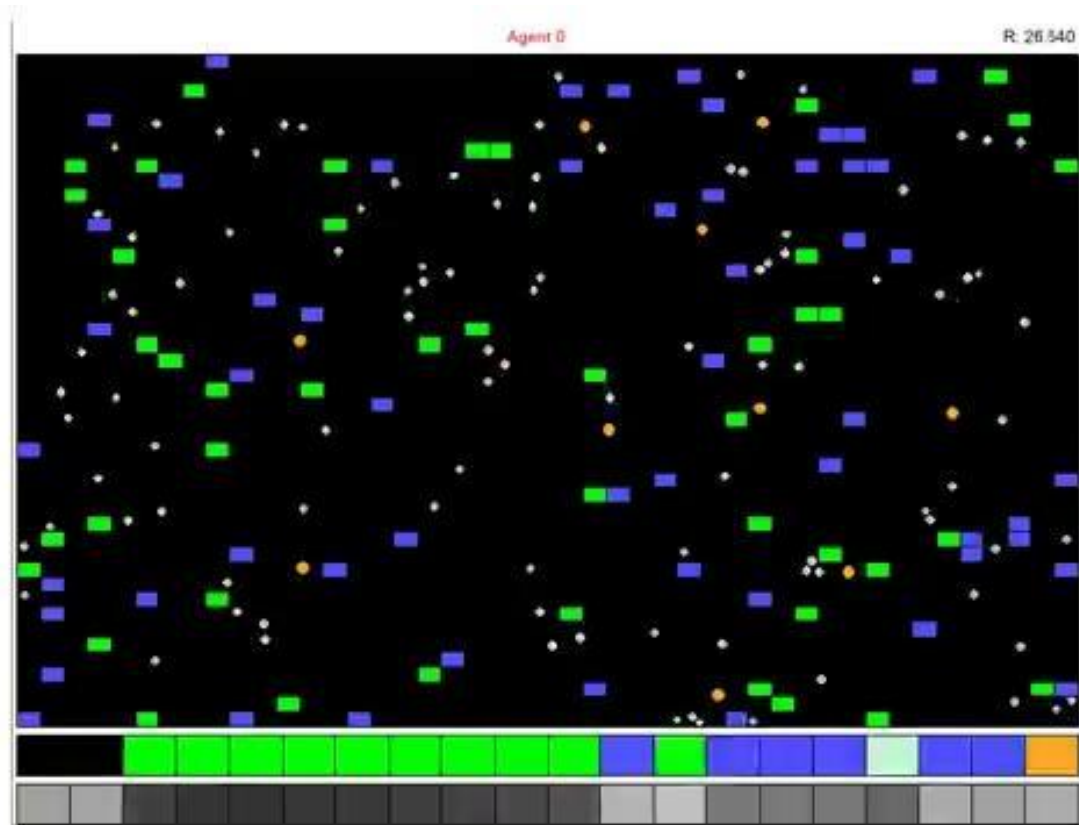


- A **bottom-up collective** system with **local** computations and feedback
- The interacting collective gives rise to a **learning system**
- Can be viewed as cultural evolution and its cumulative nature (**meme propagation**)



Memetic learning in a multi-agent scenario

- 100 agents, food and predators
- Memes are behaviors exchanged between agents
- Agents learn to forage food and evade predators using memetic evolution



Loss-Guarded L2O At ICML 2022




#tbt Our [#ICML2022](#) presentation on neural network optimization with Loss-Guarded L2O (LGL2O).

👁 Watch the video presentation here:
slideslive.com/38983756/a-sim...

📄 Paper: arxiv.org/abs/2201.12426

🌐 ICML Conference @icmlconf · Sep 5
🔗 #icml2022 video presentations are now available. 🤖 🌟
slideslive.com/icml-2022



A Simple Guard for Learned Optimizers



Jaroslav Vítků, Isabeau Prémont-Schwarz, Jan Feyereisl



LLMs - what are they?

Quick demo with OpenAI API



Would you help
me with this
barrel?



Anything for you
love.

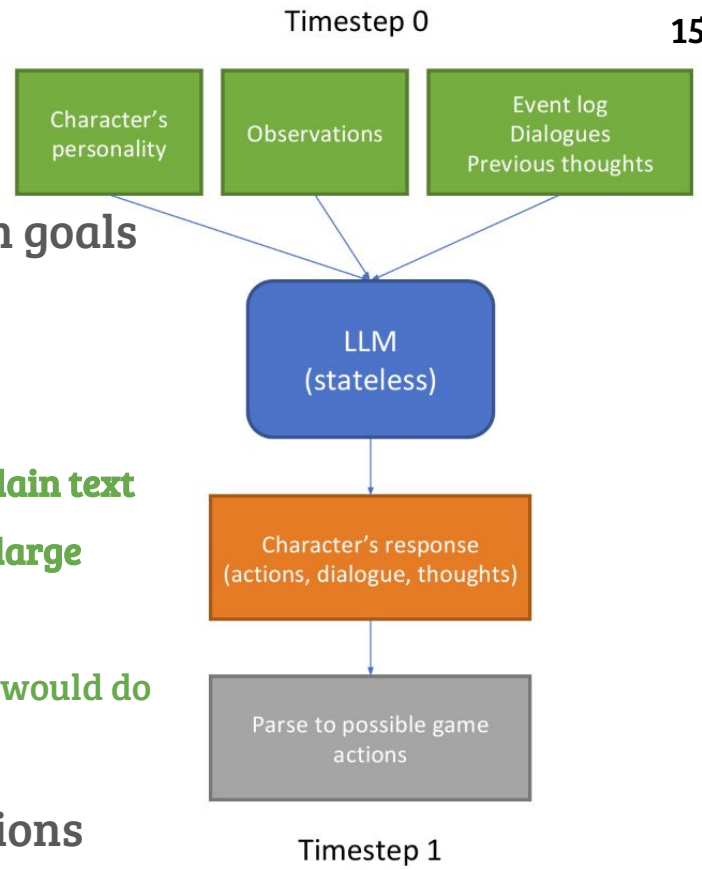


AI Game



AI Game

- AI agents are emulated personalities with goals and memories
- How does it work?
 - Agent's personality described by designers in **plain text**
 - Observations and recent events submitted to a **large language model (LLM)**
 - LLM predicts what this person in this situation would do => **emulated personality**
- Output is translated to possible game actions



AI Game

Large pre-trained language models

- Trained on internet-scale text corpus with a single objective: predict the next word (auto-complete)
- Runs on server GPUs, 20-500 GB of VRAM, inference latency 100+ milliseconds
- Inference is expensive (\$1 / 1 hour / 1 player)
- Rough estimation: 10 GPUs, 1 FPS (LLM inference) -> versus <- 1 GPU, 60 FPS (rendering)
- Examples: GPT-3, GPT-NeoX, BLOOM

What is missing in our current AI models?

- Agents hallucinate irrelevant things or absolute nonsense
- Limited long-term memory (forgets quickly, ignores things in memory)
- Language models don't improve by experience
- Observational space is limited



AI Game

Large pre-trained language models

- Trained on internet-scale text corpus with a single objective: predict the next word (auto-complete)
- Runs on server GPUs, 20-500 GB of VRAM, inference latency 100+ milliseconds
- Inference is expensive (\$1 / 1 hour / 1 player)
- Rough estimation: 10 GPUs, 1 FPS (LLM inference) -> versus <- 1 GPU, 60 FPS (rendering)
- Examples: GPT-3, GPT-NeoX, BLOOM

What is missing in our current AI models?

- Agents hallucinate irrelevant things or absolute nonsense
- Limited long-term memory (forgets quickly, ignores things in memory)
- Language models don't improve by experience
- Observational space is limited

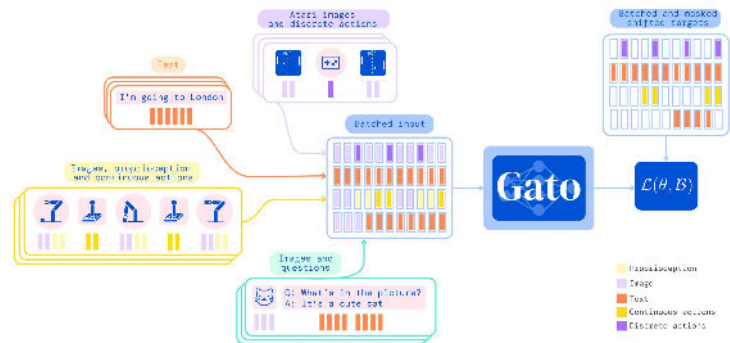
Direct research motivation!



LLM improvements

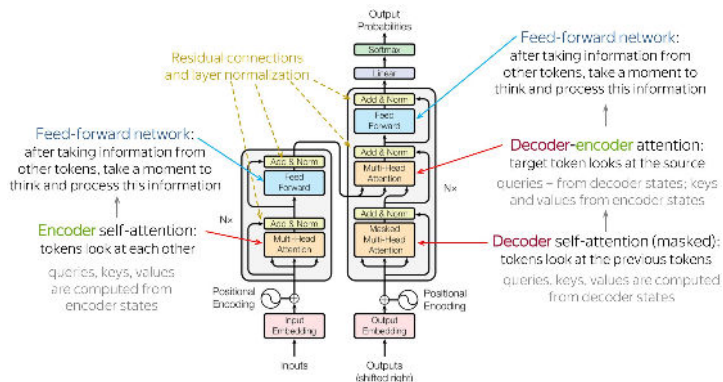
Motivations driven by AI Game

- Better respect facts about the world
- Retrieve relevant facts from storage
- Continually adapt to the player



Research topics


- Reasoning that scales with number of facts
- Growing long-term memory
- Online continual learning



Open House Day of GoodAI and Keen SWH by [Stable Diffusion](#)




Open House Day of GoodAI and Keen
Software House in Prague, 2022, oil painting



Open House Day of GoodAI and Keen
Software House in Prague, 2022, futuristic




Rate of progress: expect the unexpected



2) Teaching robots to walk in Space Engineers

Walking Robots - aren't they just cool??



Let's get walking robots to Space Engineers!



VeriDream



“VERTical Innovation in the Domain of Robotics Enabled by AI Methods”

- Connecting research and business partners for applied research
- EU-funded, international cooperation



DLR

Deutsches Zentrum
für Luft- und Raumfahrt
German Aerospace Center



INSTITUT
POLYTECHNIQUE
DE PARIS



GoodAI



MAGAZINO

“Allow the players to build, animate and control their own robots in Space Engineers”


Necessary steps

- Inverse kinematics
- Gait controller
- Robustness of the gait






Leg controller

- Control by target position ->
Inverse Kinematics



Leg controller

- Possible solutions:
 - Numerical IK (e.g. [OpenRAVE](#))
 - Optimization algorithm (e.g. [Quality Diversity](#), [arXiv:2109.06826](#))






Leg controller

- Possible solutions:



- Numerical IK (e.g. [OpenRAVE](#))
- Optimization algorithm (e.g. [Quality Diversity](#))


- + Development speed
- + Robustness
- Time to calculate
- Precision



Quality Diversity


- Evolutionary optimization algorithm
- Generates large collection of high quality diverse solutions
- Implementation available e.g. in [pyribs](#)






Gait controller

- Coordinate legs for walking
- Requirements:
 - Walk forward and backward
 - Rotate while walking/standing still
 - Control by target location or WASD
 - Support various gait styles
 - Adapt to terrain




Gait controller

- Optimization problem - 6 scalar parameters
- Optimizer goal is to walk fast in a straight line
- Quality diversity can be used again
- Another option: [Optuna](#)



Gaits found by optimizer



Terrain adaptation




Other approach - Reinforcement learning

- Robot is controlled by a neural network
- Movements are generated as a reaction to the environment
- Nice moves are rewarded, bad moves are penalized
- What is nice? What is bad?

Implemented using RL framework [Stable-Baselines3](#), [jmlr:v22/20-1364](#)




RL controller




RMA: Rapid Motor Adaptation for Legged Robots

Method used to train neural network in simulation and then use it on real world robots ([arXiv:2107.04034](https://arxiv.org/abs/2107.04034))



RMA - Isaac Gym

- Physics simulation environment for reinforcement learning
- Allows running simulations of thousands environments in parallel
- Super fast - everything is running on GPU
- Conversion from one environment to another may be tricky



[\(arXiv:2108.10470\)](https://arxiv.org/abs/2108.10470)



Takeaways from research cooperation

1. Define and constrain the goal well
2. Align partner motivations
3. Share knowledge and insights, not just code
4. Iterate quickly
5. Define proxy evaluation metrics besides KPIs
6. Collaboration at a distance requires more effort at communication



We are hiring

- **GoodAI - AI Game**
 - **Community Manager**
 - **PR Specialist**
 - **ML / NLP engineers**
- **Keen Software House**
 - **Senior Programmers**
 - **Senior Artists**
 - **UI Designers**



Questions?



VeriDream - DLR

