

FC Portugal: Tri-World Champions in RoboCup 3D Humanoid Soccer Simulation

Luís Paulo Reis

lpreis@fe.up.pt

Associate Professor at FEUP - Faculty of Engineering of the University of Porto
Director of LIACC – Artificial Intelligence and Computer Science Laboratory, Univ. Porto
Member of Coordinating Commission of LASI – Intelligent Systems Associate Laboratory
President of the GA of APPIA – Portuguese Association for Artificial Intelligence



Agenda

- **AI and Machine Learning**
- **Generative AI and LLMs**
- **AI in Research and Robotics**
- **Agentic AI, Robot Learning, DRL, LBM**s
- **RoboCup and Robotic Soccer**
- **FC Portugal Team**
- **Humanoid Skills with DRL**
- **Conclusions**

AI and the Discovery of the New World

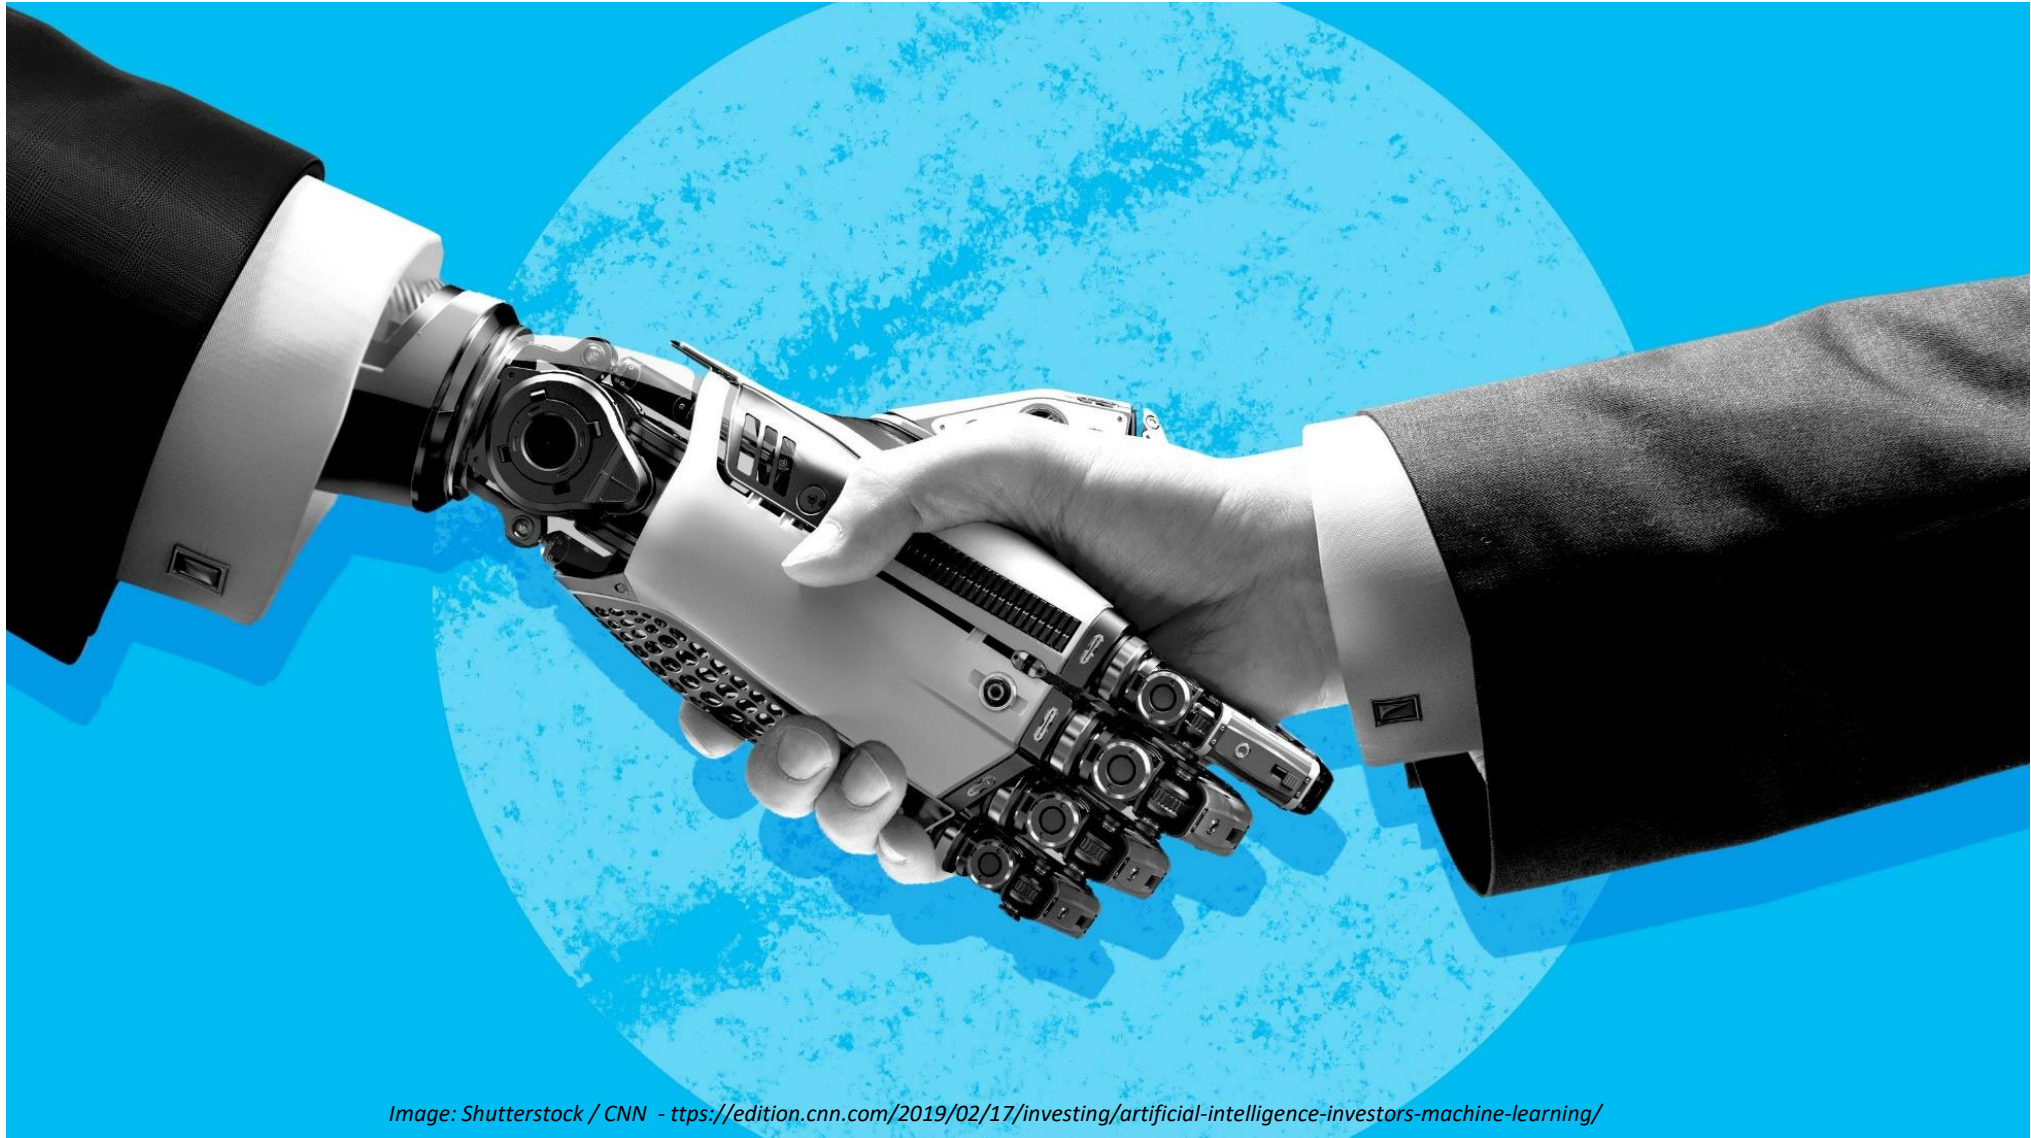


Image: Shutterstock / CNN - <https://edition.cnn.com/2019/02/17/investing/artificial-intelligence-investors-machine-learning/>

Weak and Strong AI

- **Weak Artificial Intelligence**

weak AI, also known as narrow AI is artificial intelligence that is focused on one **single narrow task**

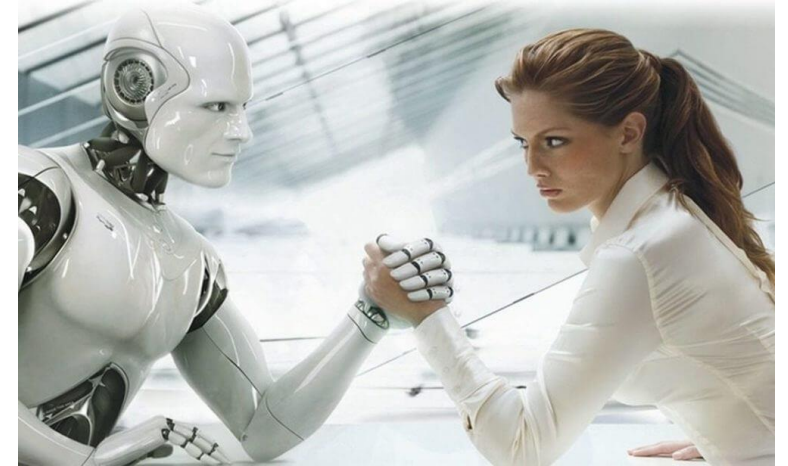


Image: <https://livingsmartheart.com/what-is-strong-ai/>

- **Strong Artificial Intelligence**

Strong AI or Artificial General Intelligence (AGI) is the intelligence of a machine that could successfully perform **any intellectual task** that a human being can!

Science fiction?



Machine Learning

Field of artificial intelligence that gives **computer systems** the **ability to "learn"** (e.g., progressively **improve performance** on a specific task) from **data/results of their actions**, without being explicitly programmed

Supervised Learning

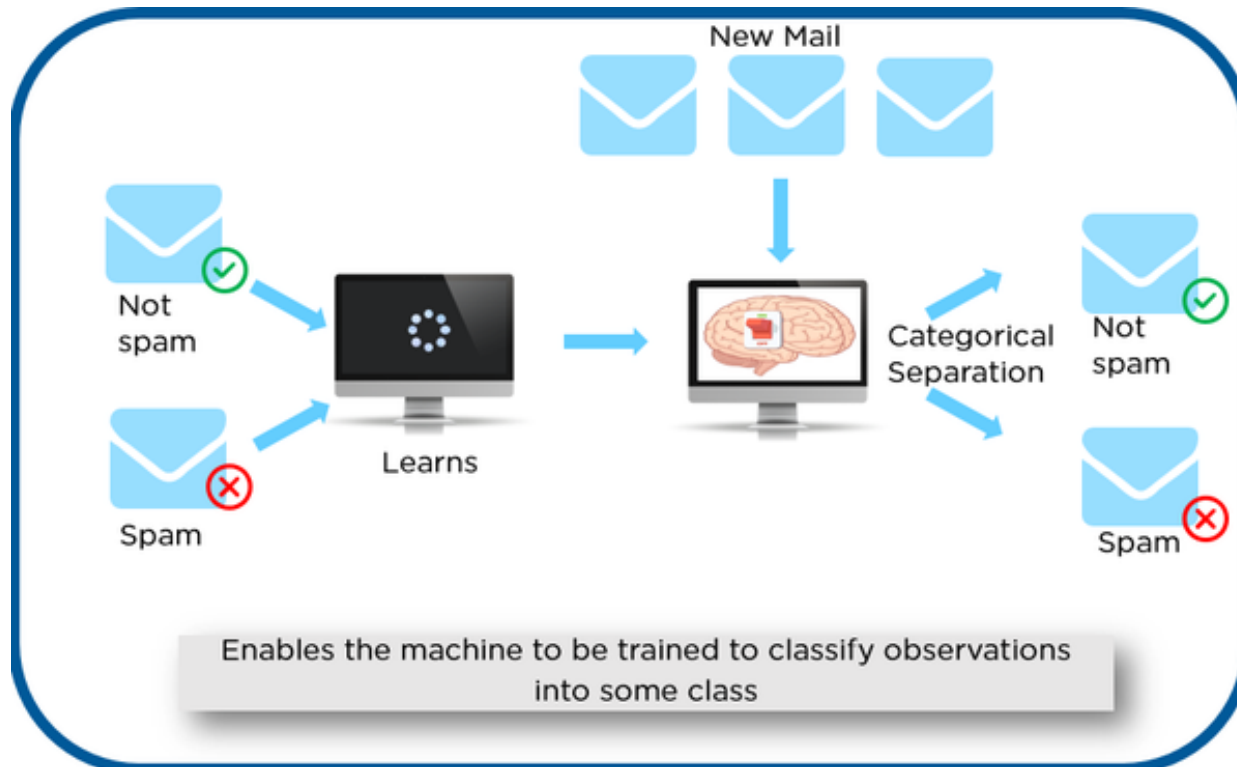


Image: <https://towardsdatascience.com/what-are-the-types-of-machine-learning-e2b9e5d1756f>

Reinforcement Learning

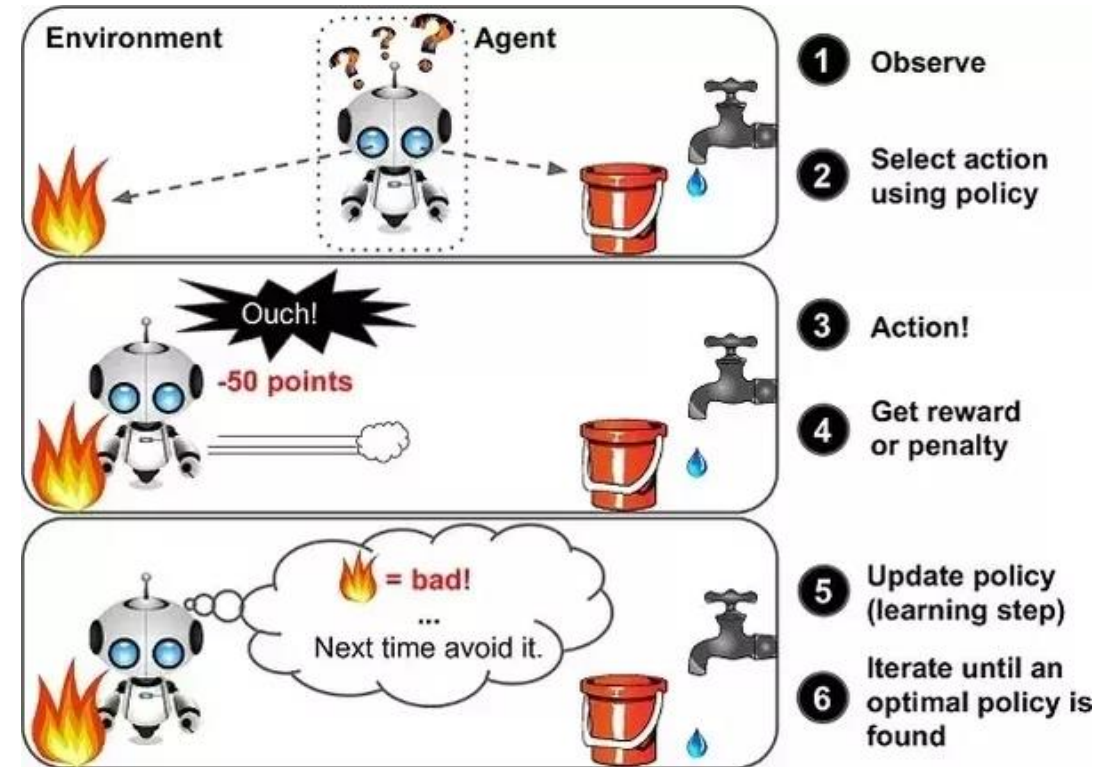
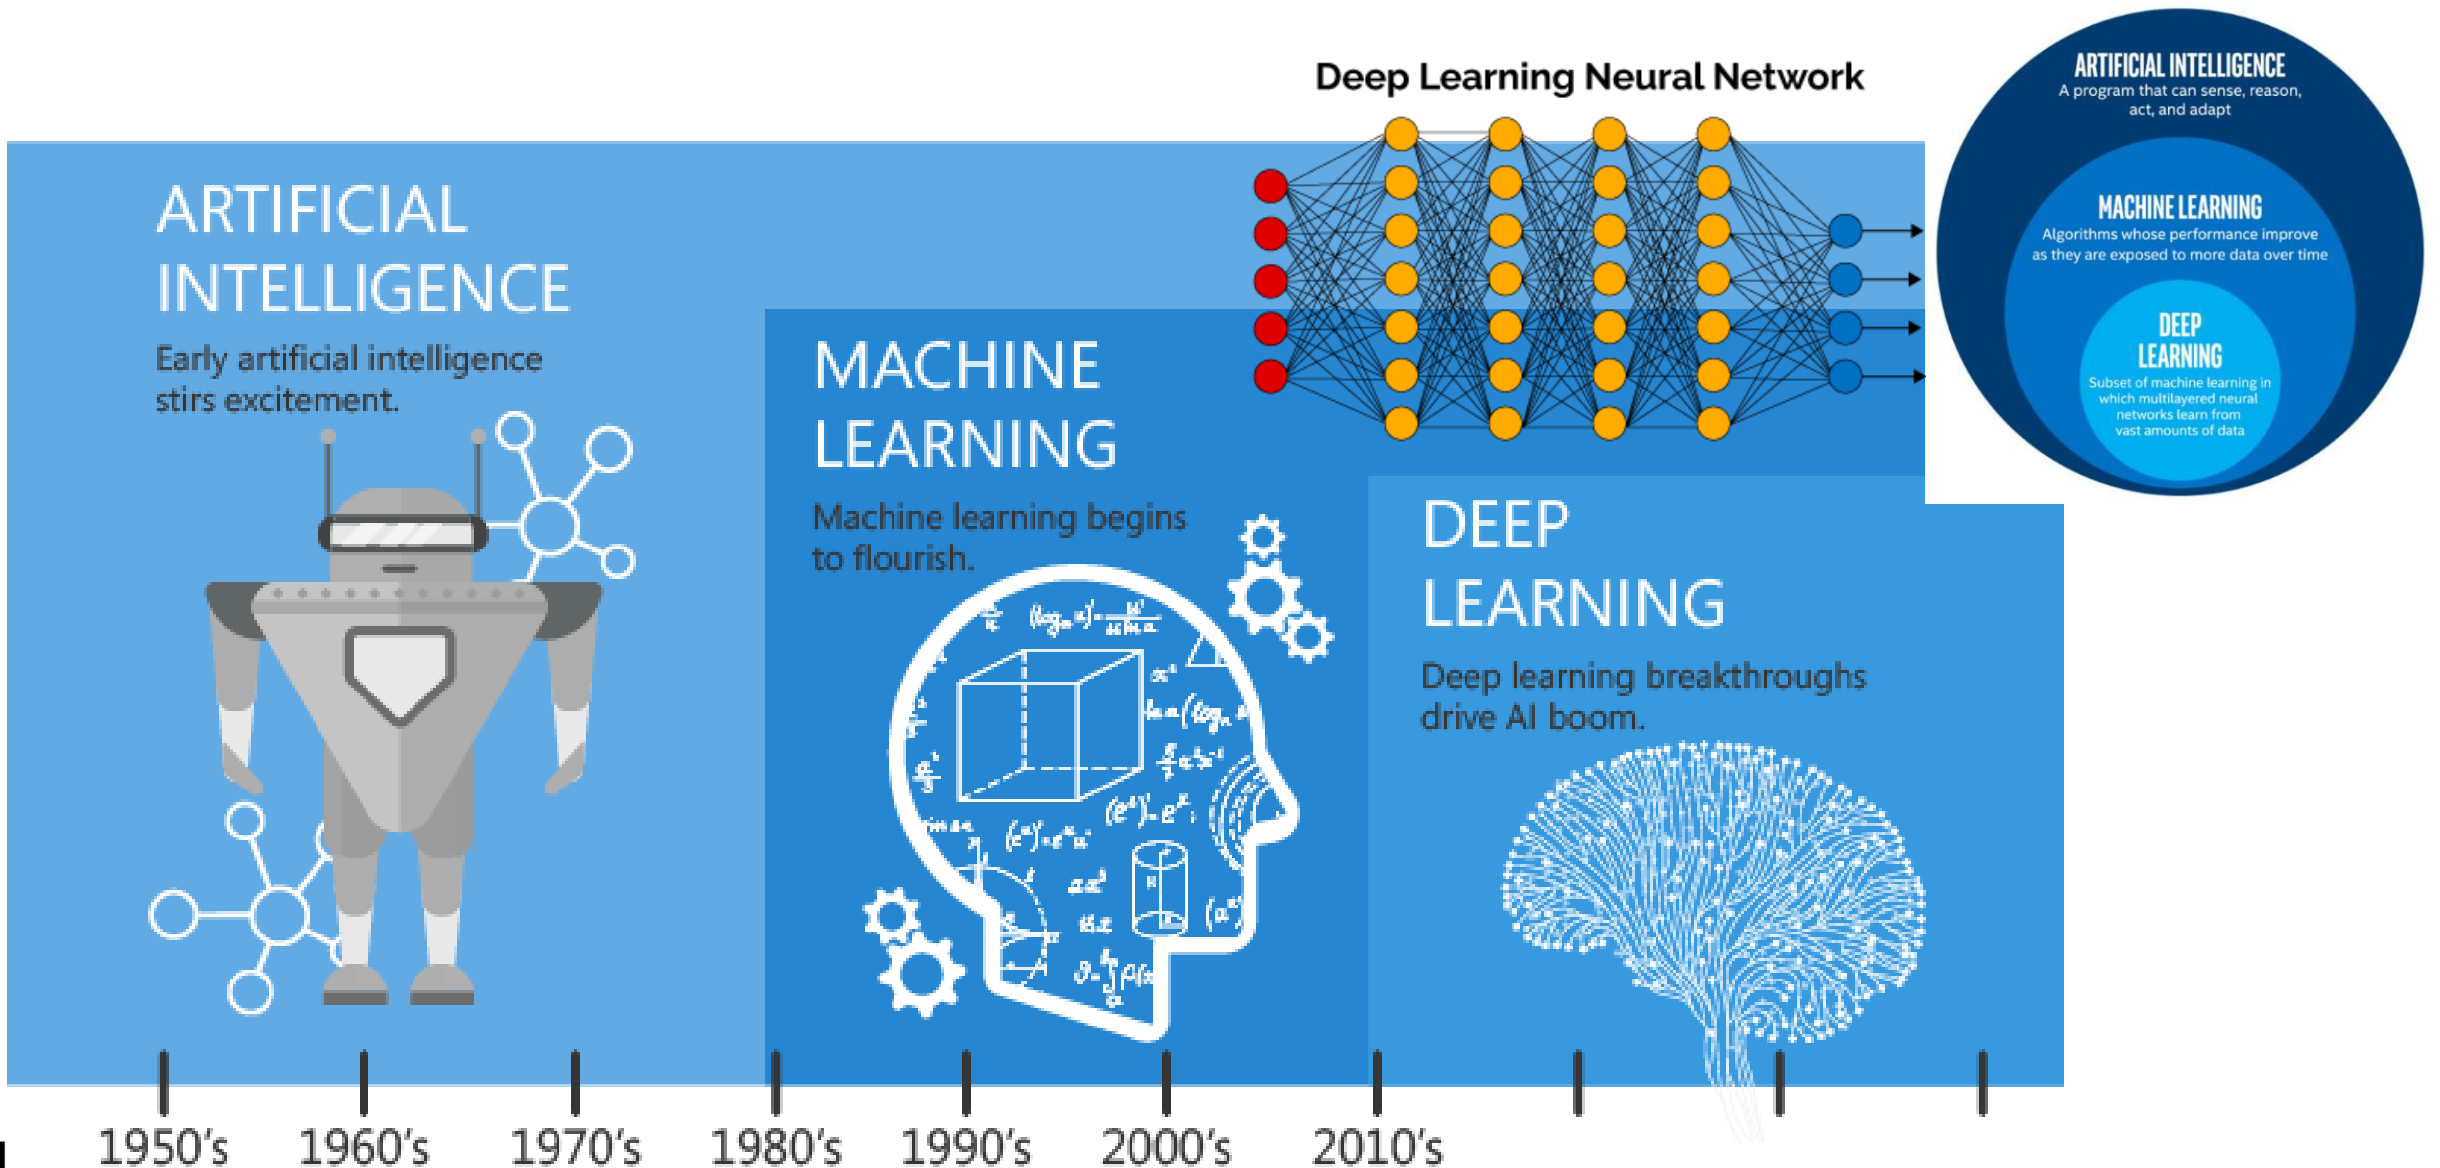


Image: <https://becominghuman.ai/an-introduction-to-machine-learning-33a1b5d3a560>

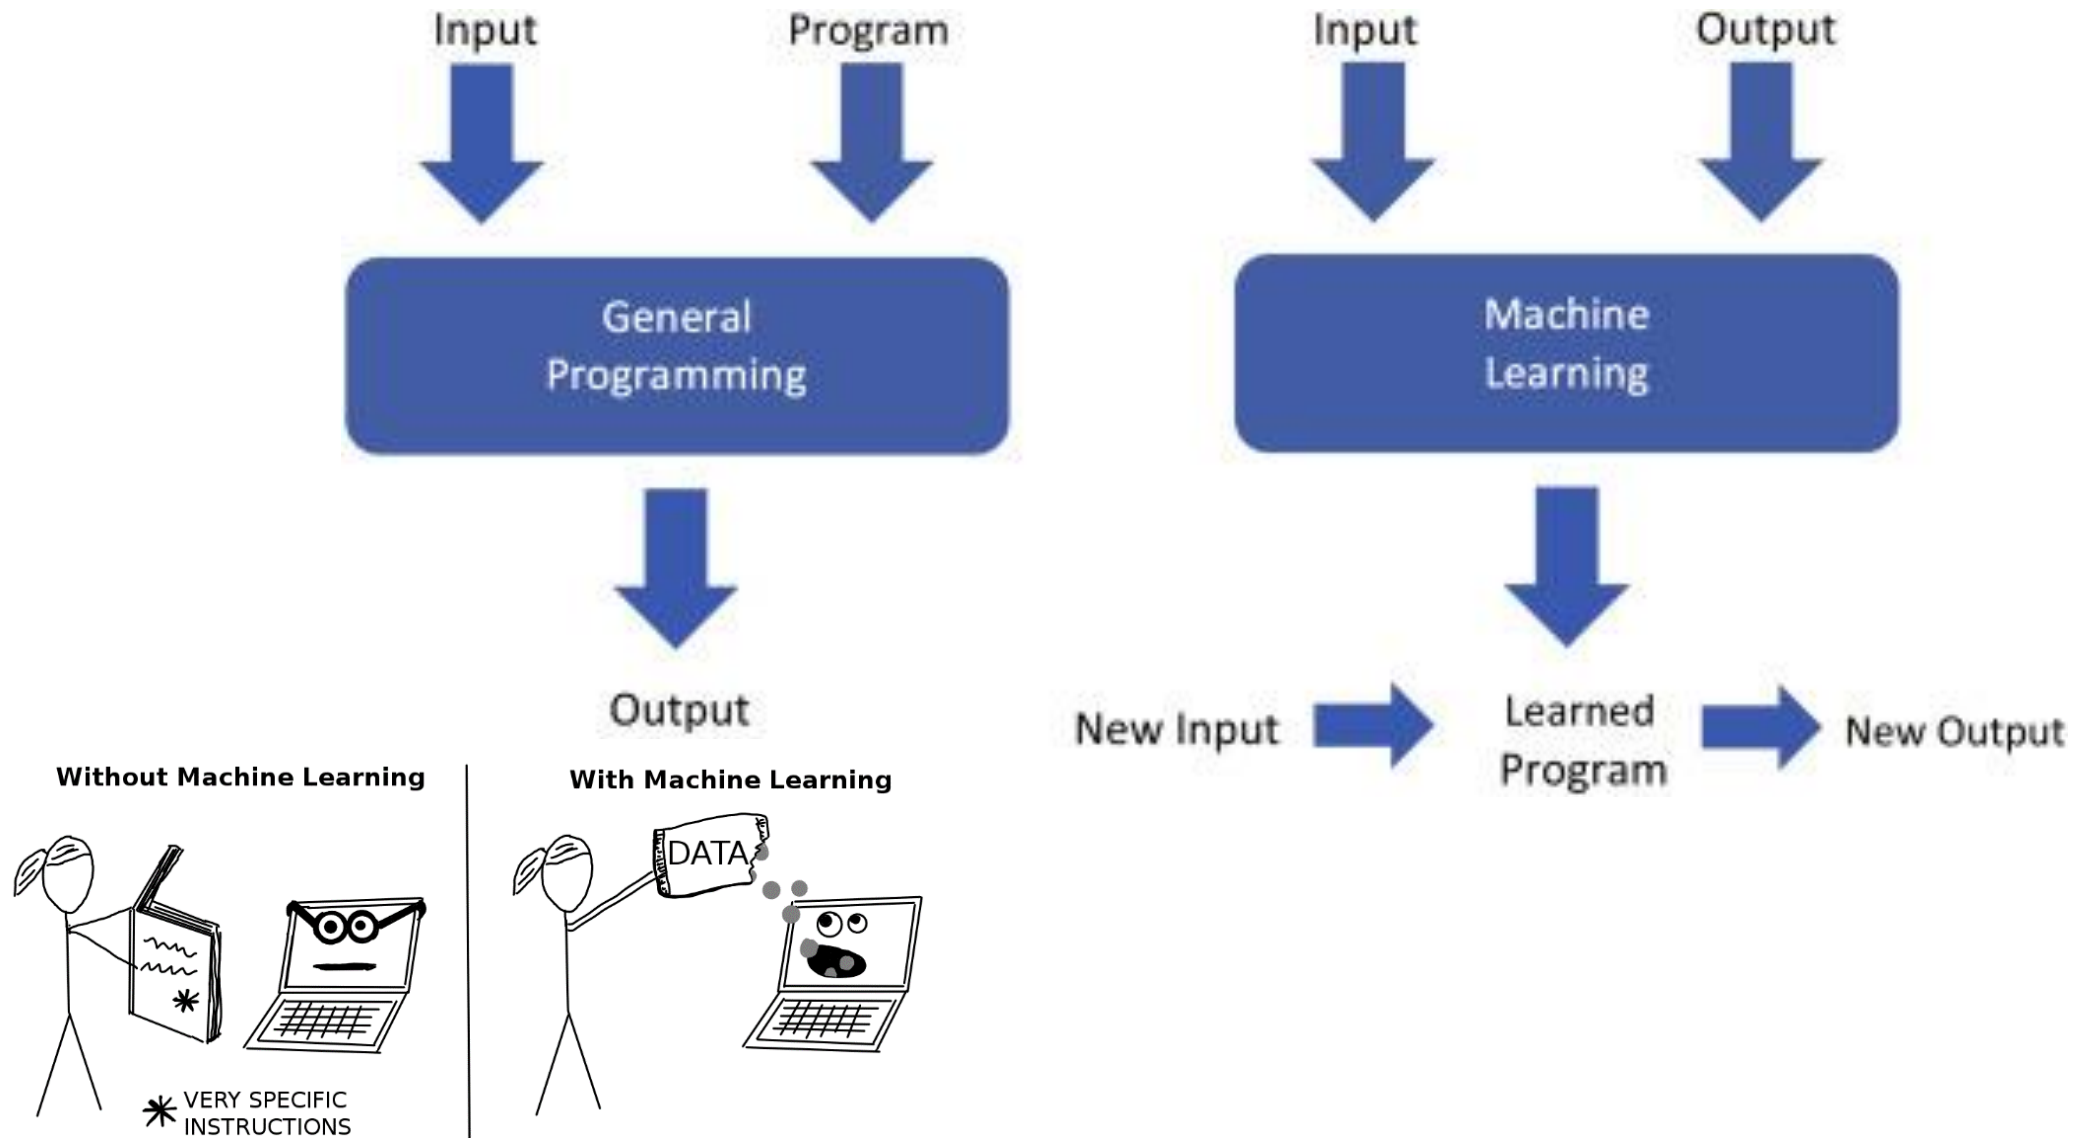
Machine Learning - History



Artificial Intelligence Today



Programming vs Machine Learning



The End of Programming!

Nvidia CEO predicts the death of coding — Jensen Huang says AI will do the work, so kids don't need to learn

News

By Benedict Collins published February 26, 2024

Jensen Huang believes coding languages are a thing of the past



(Image credit: Nvidia)

Nvidia CEO Jensen Huang has once again announced the death of coding, but this time in front of a potentially far more influential audience.

StarCoder 2 is a code-generating AI that runs on most GPUs

Kyle Wiggers @kyle_i_wiggers / 2:00 PM UTC • February 28, 2024

Comment



Image Credits: Tippiatt / Getty Images

Developers are adopting AI-powered code generators — services like [GitHub Copilot](#) and [Amazon CodeWhisperer](#), along with open access models such as Meta's [Code Llama](#) — at an [astonishing](#) rate. But the tools are far from ideal. Many aren't free. Others are, but only under licenses that preclude them from being used in common commercial contexts.

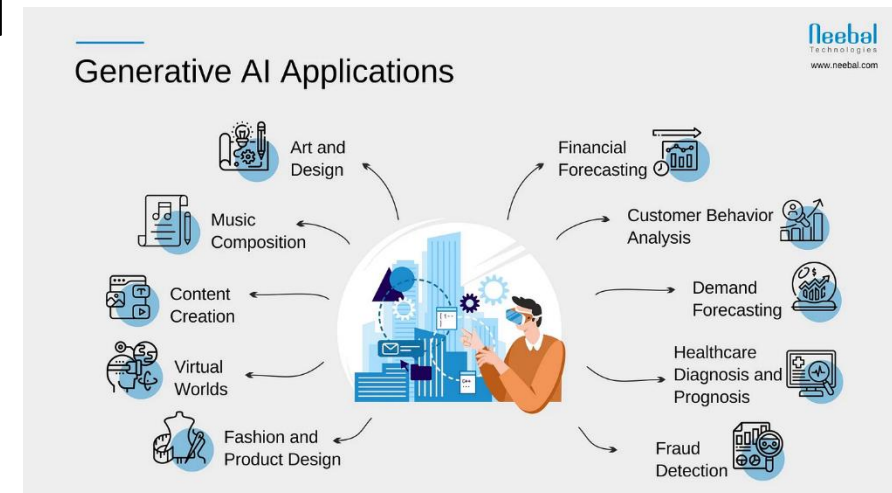
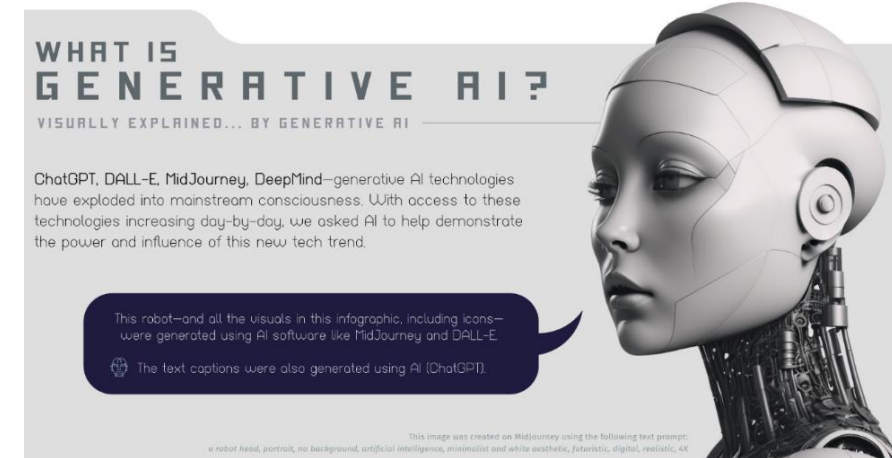
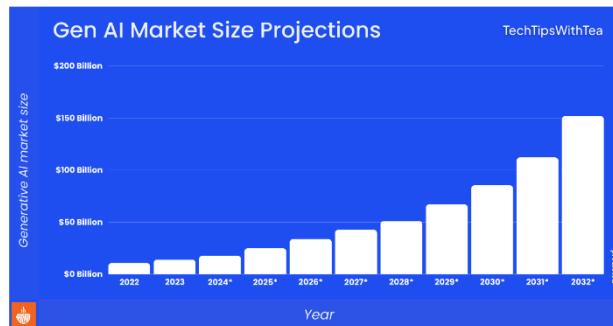
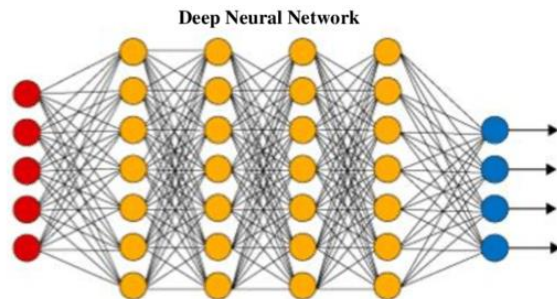
Perceiving the demand for alternatives, AI startup Hugging Face several years ago teamed up with ServiceNow, the workflow automation platform, to create [StarCoder](#), an open source code generator with a less restrictive license than some of the others out there. The original came online early last year, and work has been underway on a follow-up, StarCoder 2, ever since.

StarCoder 2 isn't a single code-generating model, but rather a family. Released today, it comes in three variants, the first two of which can run on most modern consumer GPUs:

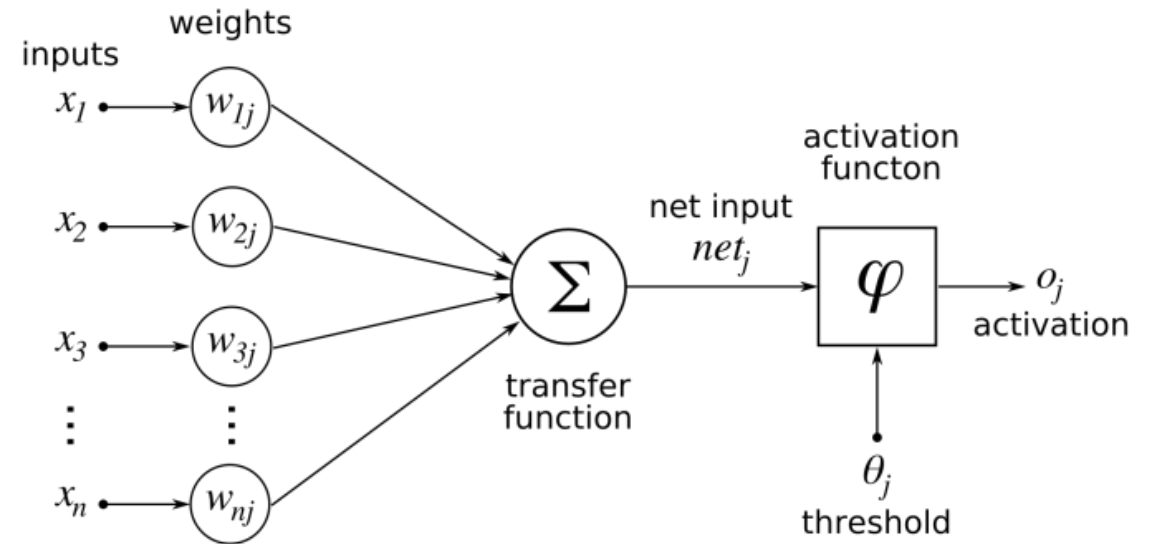
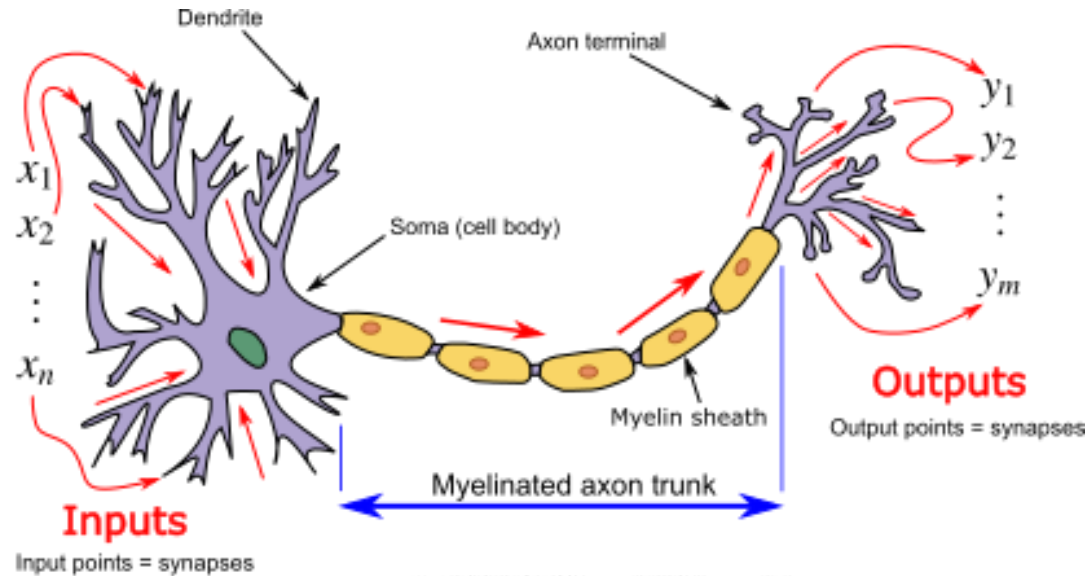
- A 3-billion-parameter (3B) model trained by ServiceNow
- A 7-billion-parameter (7B) model trained by Hugging Face
- A 15-billion-parameter (15B) model trained by Nvidia, the newest supporter of the StarCoder project

Generative AI

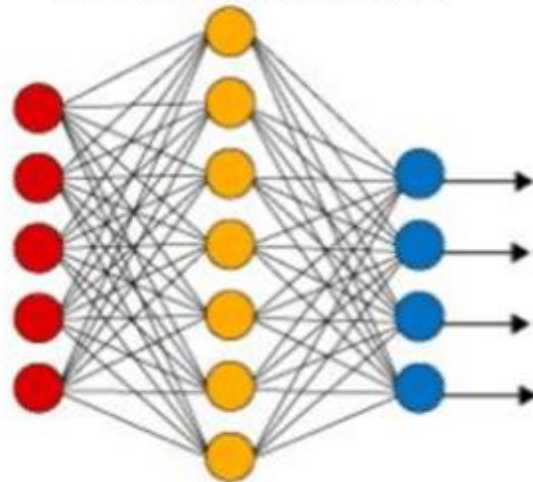
- **Generative AI (GenAI)** can create a wide variety of data, such as images, videos, audio, text, and 3D models
- **Learning patterns from existing data**, and then **using this knowledge to generate new and unique outputs in response to prompts**
- GenAI produce **highly realistic and complex content** that mimics human creativity
- **Large language models (LLM)** consist of a **very large neural network** with billions of weights, trained on very large quantities of text



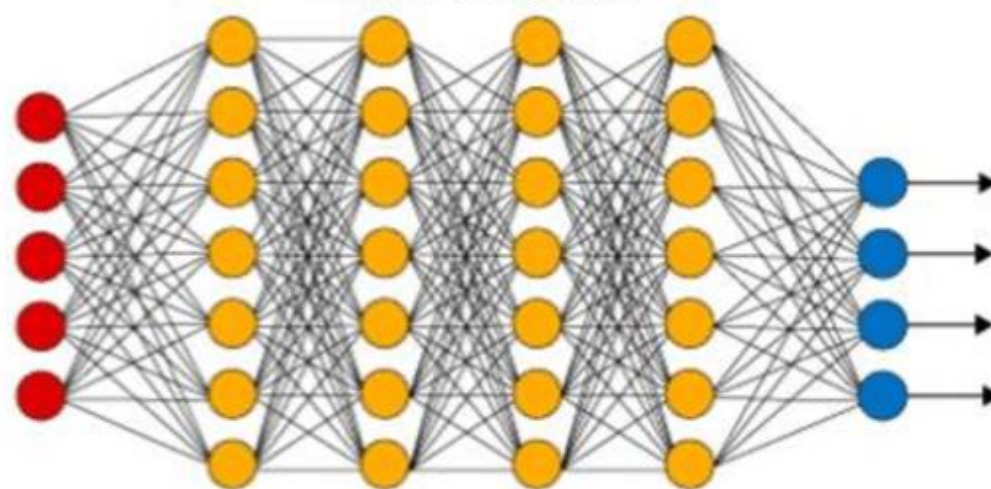
Artificial Neural Networks



Artificial Neural Network



Deep Neural Network



Generative AI Competition



LLM Evolution

2022/3->2024/5

| 2022-23 ¹ | | Jan 2025 ² | |
|----------------------|---|-----------------------|--|
| Anthropic | Claude <ul style="list-style-type: none">• Not multimodal (text only)• Limited contextual understanding (difficulty with complex conversations)• No tool usage | → | Claude 3.5 <ul style="list-style-type: none">• Multimodal (text, audio, and images)• Enhanced contextual understanding and coherence during long interactions• Experimental computer usage capability for some users |
| Google | Google Bard <ul style="list-style-type: none">• Not multimodal (text only)• Fair reasoning (eg, scored high on SAT, but bottom 10% on bar examination)• Limited contextual understanding (difficulty with coherence in complex conversations)• Limited real-time data integration• Low personalization (limited adaptability) | → | Gemini 2.0 Flash <ul style="list-style-type: none">• Multimodal (text, audio, and images)• Advanced reasoning (capable of multistep problem-solving and nuanced analysis)• Enhanced contextual understanding (maintains coherence in long dialogues)• Real-time data integration (from Google Search)• Advanced personalization (user context) |
| OpenAI | GPT-3.5 <ul style="list-style-type: none">• Not multimodal (text only)• Fair reasoning ability (eg, scored high on SAT, but bottom 10% on bar examination)• Limited contextual understanding (difficulty with coherence in complex conversations)• Standard API access (for text generation) | → | OpenAI o1 <ul style="list-style-type: none">• Multimodal (text and images)• Advanced reasoning (eg, top 10% on bar examination)• Enhanced contextual understanding (maintains coherence in long dialogues)• Advanced API access (supports multimodal inputs) |

Scopus AI



ELSEVIER

Publish with us



[Home](#) > [Products](#) > [Scopus](#) > [Scopus AI](#)

Scopus AI: Trusted content. Powered by responsible AI.

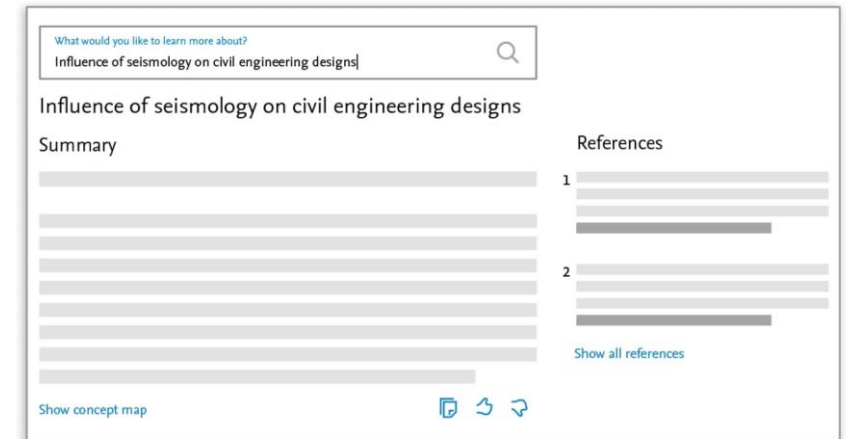
Empower your research journey with Scopus AI – your dynamic GenAI-powered research companion.

Navigate through the vast expanse of human knowledge faster with a trusted guide designed to enhance your understanding, enrich your insights, and transform your overall research experience.

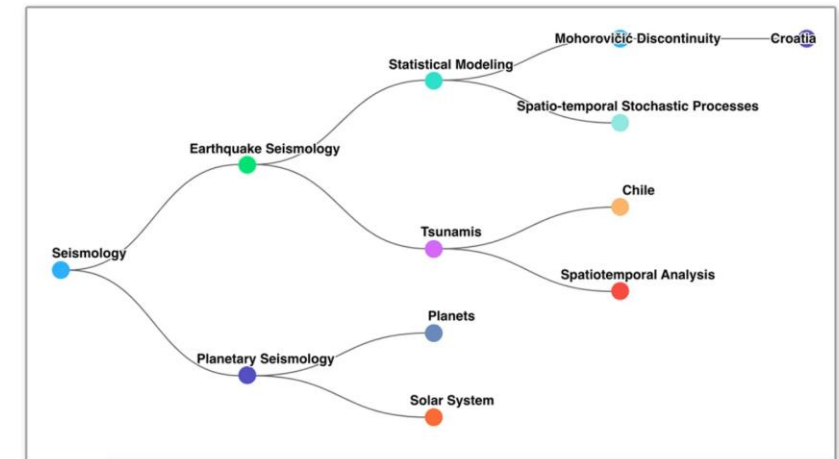


Elsevier's Five Responsible AI Principles

- We consider the real-world impact of our solutions on people
- We take action to prevent the creation or reinforcement of unfair bias
- We can explain how our solutions work
- We create accountability through human oversight
- We respect privacy and champion robust data governance



Topic summaries
— created in seconds



A more complete picture
— from our Concept map

Elicit

FAQ

Tasks

Starred



Ask a research question

Elicit will find answers from 175 million papers



Or run Elicit over your own papers

Upload PDFs

Welcome to Elicit, your AI research assistant

Select a workflow



Extract information from PDFs

Upload up to 100 papers and extract information into a structured table



Discover concepts across papers

Get a table of concepts synthesized from the research literature



Find scientific research papers

Get a table of published research papers related to your question or topic

Developing Soccer Skills in Humanoid Robots Using Deep Reinforcement Learning

Share

how to develop robust humanoid robot skills with deep reinforcement learning for soccer play...

Summary of top 4 papers

Copy

Recent research demonstrates the potential of deep reinforcement learning (Deep RL) for developing robust and dynamic skills in humanoid soccer robots. Studies have shown that Deep RL can enable low-cost humanoid robots to acquire complex movements like rapid fall recovery, walking, turning, and kicking, which can be seamlessly integrated into game strategies (Animoni [Nagaraju et al., 2023](#); Tuomas [Haarnoja et al., 2023](#)). These learned skills significantly outperform scripted baselines in speed and efficiency. Researchers have also successfully applied Deep RL to teach robots obstacle-aware dribbling ([Kuwano et al., 2022](#)) and natural running gaits ([Abreu et al., 2019](#)). The learned policies can be transferred from simulation to real robots, although some challenges remain in achieving stability for all skills in competitive environments. Key factors enabling successful transfer include high-frequency control, targeted dynamics randomization, and training with perturbations. These advancements contribute to more human-like behaviors in robot soccer and hold promise for complex multi-robot interactions.

Consensus



You

write me a state of the art about AI ethics



✓ Talked to chat.consensus.app ▾

The state of the art in AI ethics encompasses a wide range of topics, from algorithmic injustice and discrimination to labor impacts, privacy concerns, and the societal implications of AI. Here's a summary of key findings from recent research:

1. **Algorithmic Injustice and Discrimination:** AI technologies often reflect and perpetuate existing societal biases, leading to concerns about algorithmic injustice and discrimination. This issue is especially pertinent in the context of societal and labor impacts, as well as in the development of ethical AI frameworks ([Gupta et al., 2021](#)).
2. **Ethical AI Development:** There is a growing emphasis on the need to develop AI systems ethically. This includes addressing issues of privacy, security, and the risk of misinformation. The role of AI in societal contexts, such as labor and social media, necessitates careful ethical consideration ([Vakkuri et al., 2020](#))



3. **Challenges in Implementing AI Ethics:** Despite the proliferation of ethical guidelines for AI, there's a notable gap between these guidelines and their implementation in practice. This gap is evident in industries and companies dealing with AI, indicating a need for better integration of ethical principles in AI development ([Vakkuri et al., 2022](#)).
4. **AI and Labor Impact:** The impact of AI on labor and employment is a significant ethical concern. AI's potential to automate jobs raises questions about job displacement and the need for new skills and training for the workforce ([Gupta et al., 2021](#)).
5. **Privacy and Security:** With the increasing use of AI in various domains, issues of data privacy and security are paramount. Ensuring the protection of personal data and preventing misuse of AI technologies are key ethical challenges ([Gupta et al., 2021](#)).
6. **Ethical Implications in Specific AI Applications:** Specialized areas like AI-generated art and machine translation also present unique ethical challenges. For instance, AI-generated art raises questions about the representation of artistic styles and cultural nuances ([Srinivasan & Uchino, 2021](#)).



Scite AI

scite_

Assistant

Product ▾

Solutions ▾

Blog



give me a state of the art about reinforcement learning for humanoid robotic soc...

[Export References](#)

↳ give me a state of the art about reinforcement learning for humanoid robotic soccer

2 free prompt(s) left. [Start your free trial](#) to get unlimited prompts.



Reinforcement learning (RL) has emerged as a pivotal technique in the development of humanoid robots for soccer applications, enabling these robots to learn complex behaviors and strategies in dynamic environments. This state-of-the-art overview synthesizes recent advancements in RL methodologies applied to humanoid robotic soccer, highlighting key contributions and ongoing challenges.

One of the significant areas of research involves the integration of active vision systems with RL to enhance the decision-making capabilities of humanoid soccer robots. Khatibi et al. propose a deep RL approach that optimizes the robot's viewpoint to maintain the ball within its field of vision while

↳ Ask a question... (type '/' for menu)

complemented by Haarnoja's work, which demonstrates the

References

Search Strategy



"...In this paper, we present an active vision method using a deep reinforcement learning approach for a humanoid soccer-playing robot. The proposed method adaptively optimises the viewpoint of the robot to acquire the most useful landmarks for self-localisation while keeping the ball into its viewpoint. Active vision is critical for humanoid decision-maker robots with a limited field of view..." [>>](#)

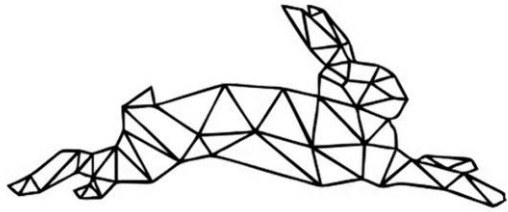
Section: Abstract

Real-time Active Vision for a Humanoid Soccer Robot Using Deep Reinforcement Learning

[Soheil Khatibi](#)¹, [Meisam Teimouri](#)², [Mahdi Rezaei](#)³ 2020 Preprint



Research Rabbit



New Collection

New Category

Connect to Zotero

Uncategorized

+ Collection

Untitled Collection

Filter

Custom

Abstracts

Comments

Select None

Select All

Untitled Collection

Kasaei

...

Li

2023

Learning hybrid locomotion skills—Learn to exploit residual actions and modulate model-based gait control

Frontiers in Robotics and AI

16 Selected Papers

Add to:

Untitled Collection

Add to Other Collection

EXPLORE PAPERS

Similar Work

Earlier Work

Later Work

EXPLORE PEOPLE

These Authors

Suggested Authors

EXPLORE OTHER CONTENT

Linked Content

EXPORT PAPERS

BibTeX

RIS

CSV

1 selected paper

Mohammadreza Kasaei

...

Zhibin Li

Learning hybrid locomotion skills—Learn to exploit residual actions and modulate model-based gait control

Frontiers in Robotics and AI

2023

PDF

This work has developed a hybrid framework that combines machine learning and control approaches for legged robots to achieve new capabilities of balancing against external perturbations. The framework embeds a kernel which is a model-based, full parametric closed-loop and analytical controller as the gait pattern generator. On top of that, a neural network with symmetric partial data augmentation learns to automatically adjust the parameters for the gait kernel, and also generate compensatory actions for all joints, thus significantly augmenting the stability under unexpected perturbations. Seven Neural Network policies with different configurations were optimized to validate the effectiveness and the combined use of the modulation of the kernel parameters and the compensation for the arms and legs using residual actions. The results validated that modulating kernel parameters alongside the residual actions have improved the

Remove from:

Untitled Collection

Add to Other Collection

EXPLORE PAPERS

Similar Work

All References

All Citations

EXPLORE PEOPLE

These Authors

Suggested Authors

EXPLORE OTHER CONTENT

Linked Content

EXPORT PAPERS

BibTeX

RIS

CSV

PUBLIC COLLECTION

SHAREABLE LINK

Copy

All References

Filter

Relevance

Abstracts

Comments

Select None

Select All

Handžić

...

Reed

2015

Perception of gait patterns that deviate from normal and symmetric biped locomotion.

Frontiers in Psychology

Evans

...

Batchelder

2012

Human preferences for symmetry: subjective experience, cognitive conflict and cortical brain activity.

PLOS ONE

Englsberger

...

Albu-Schaffer

2015

Three-Dimensional Bipedal Walking Control Based on Divergent Component of Motion

IEEE Transactions on Robotics

Robinson

...

Nigg

1987

Use of force platform variables to

Connections

Click to Hide

Connections between your collection and 16 papers

Graph Type

Network

Timeline

Labels

First Author

Last Author

Filter these items

Kasaei

2023

Franco-Robles

2020

Krishna

2021

Koryakovskiy

2018

Englsberger

2015

Handžić

2015

Caron

2020

Abdolhosseini

2019

Robinson

1987

Lin

2020

Yano

2020

Chatzinkolaides

2020

Evans

2012

Kasaei

2021

Farali

2019

Posa

2017

Ahn

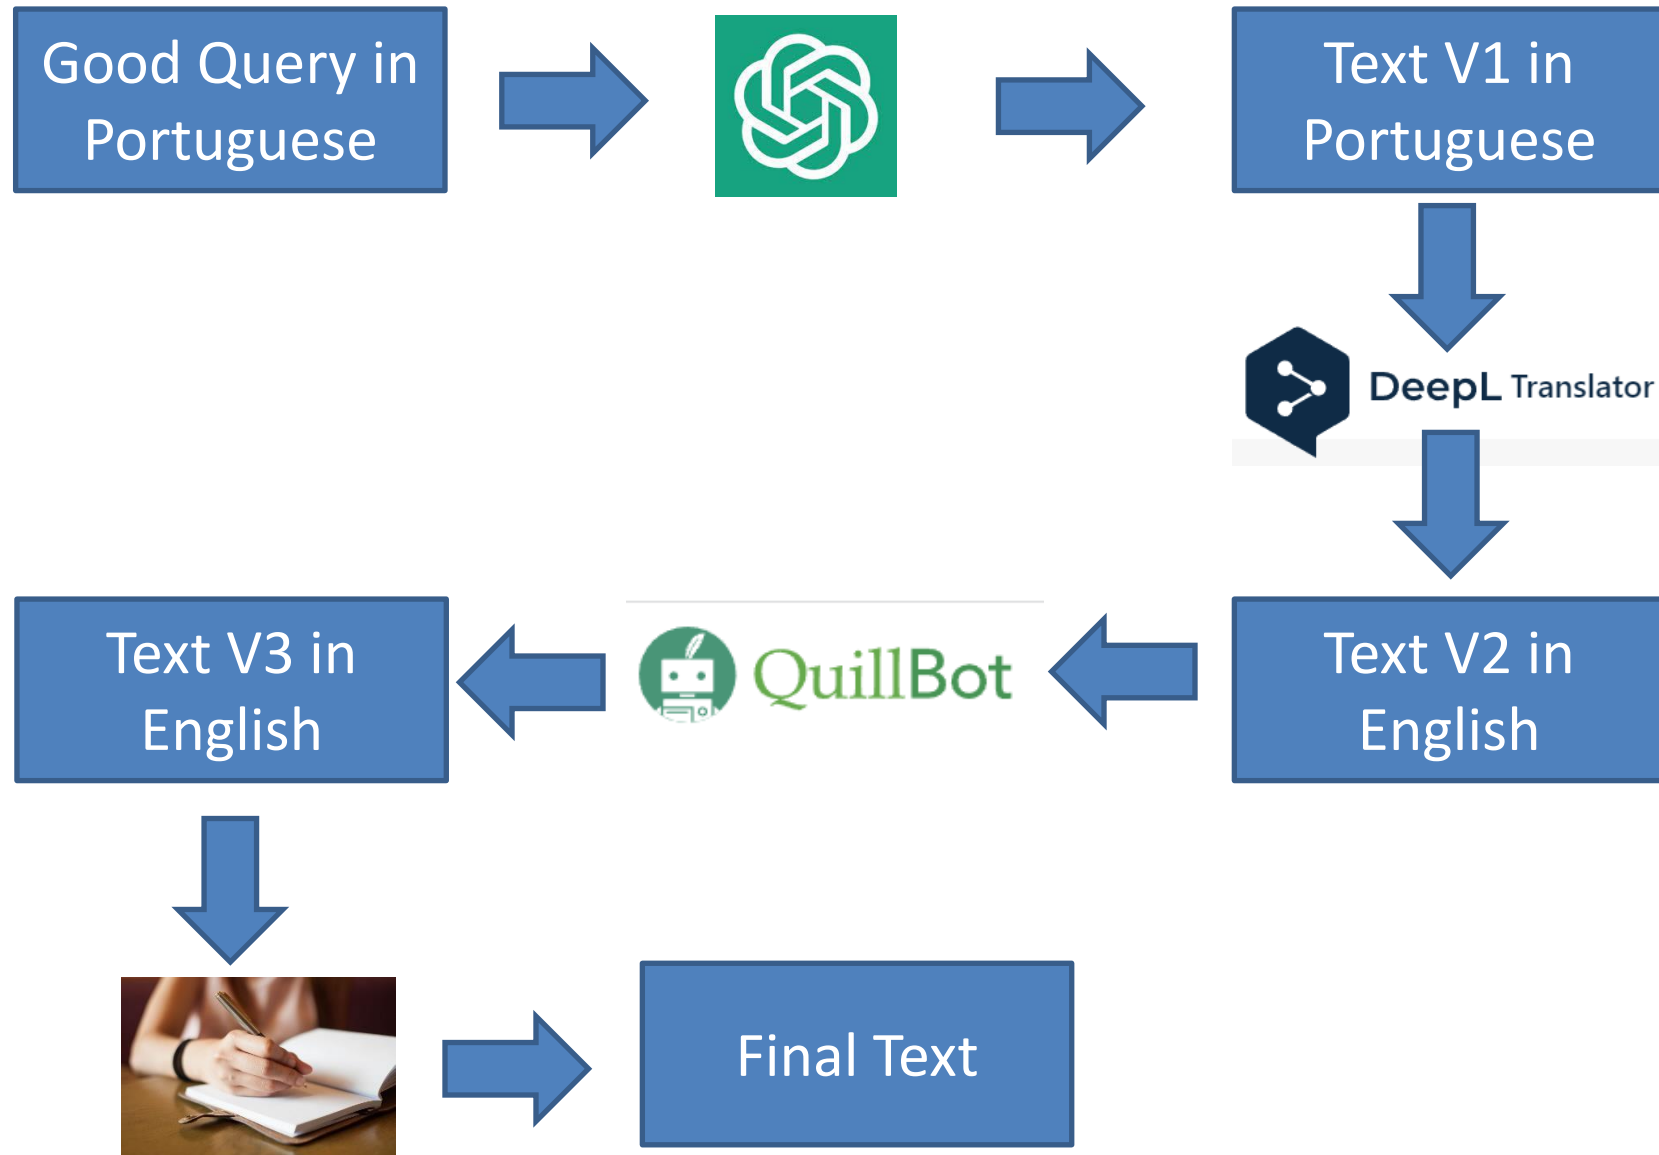
2020

Zoom Out

Fit All

Zoom In

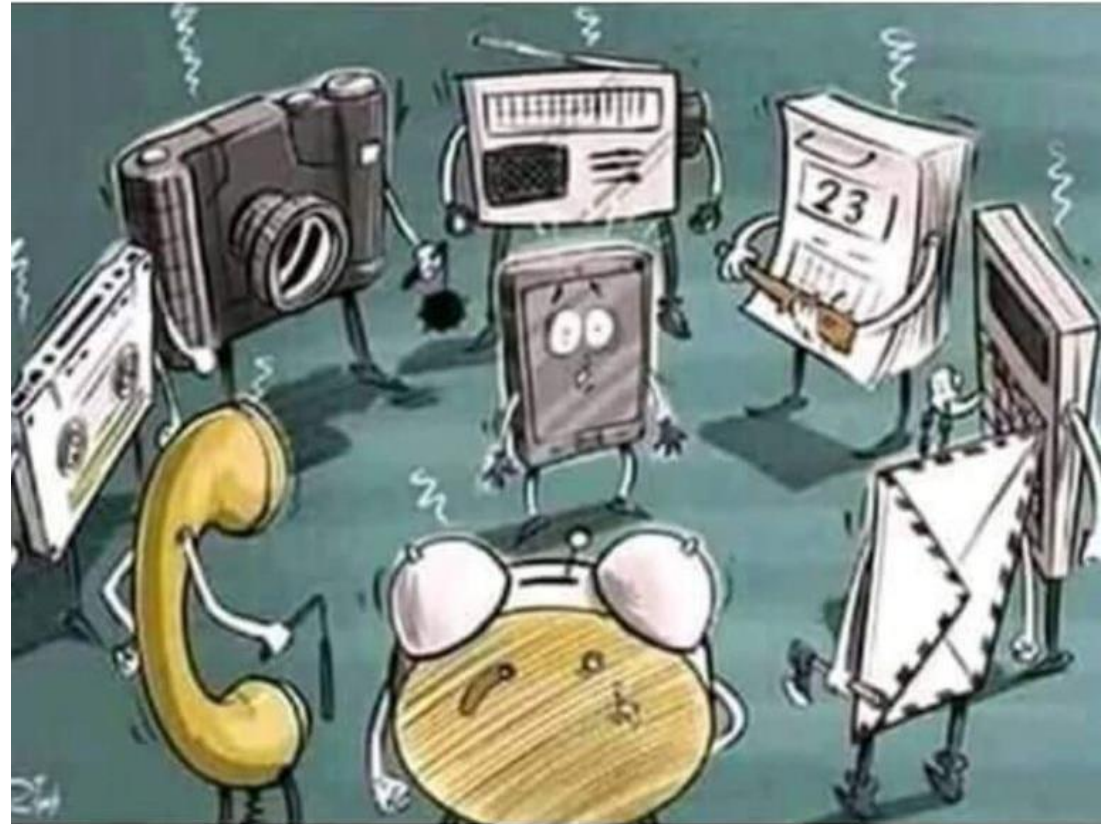
How to Use AI and Anti-Anti ChatGPT



Technology: Adapt or Resist?!

So! You are the guy...

- **COVID-19 - Use of technology**
- **Educate New Generations**
- **Ethical Use**
- **Technical Use**
- **AI Policy**
- **Prompt Engineering**



That took all of our Jobs!

Llama 3 - Meta



Technology ▾

Get started ▾

Trust & Safety ▾

Community ▾

Try Meta AI ↗

Download models



Build the future of AI with Meta Llama 3

Now available with both 8B and 70B pretrained and instruction-tuned versions to support a wide range of applications

Get Started

Experience Llama 3 on Meta AI




Mistral


Frontier AI in your hands

Open and portable generative AI for devs and businesses.


[Try the Chat](#) [Build on la Plateforme](#)

**Openness**


We lead the market of open source generative technologies to bring trust and transparency in the field and foster decentralised technology development

**Portability**

Our technology is available through serverless APIs, public cloud services (on Azure AI, on Amazon Bedrock), and for VPC/on-premise deployment. Our independence from cloud providers is a guarantee of our customers' independence.

**Value and speed**

Our flagship model, Mistral Large, has independently validated top-tier reasoning in multiple languages. All our models bring unmatched value and latency at their price points.

**Customisation**

Our models can be fine-tuned and modified at will for your business to create differentiated AI applications.

Build on Open Source

Under the Apache 2.0 license, our 3 open source models Mistral 7B, Mixtral 8x7B, Mixtral 8x22B are usable and customisable for a variety of use cases. They can be downloaded or used on demand via our platform.

- [Download them](#) for deployment in your own environment
- Use them on [La Plateforme](#) at market-leading availability, speed, and quality control



Gemma 2 – Gemini Google

Gemma

google/gemma

Gemma 2 is now available to researchers and developers

Gemma 2 offers best-in-class performance, runs at incredible speed across different hardware and easily integrates with other AI tools.

Jun 27, 2024 · 4 min read



Gemma is a family of lightweight, open models built from the research and technology that Google used to create the Gemini models.

Responsible Generative AI Toolkit

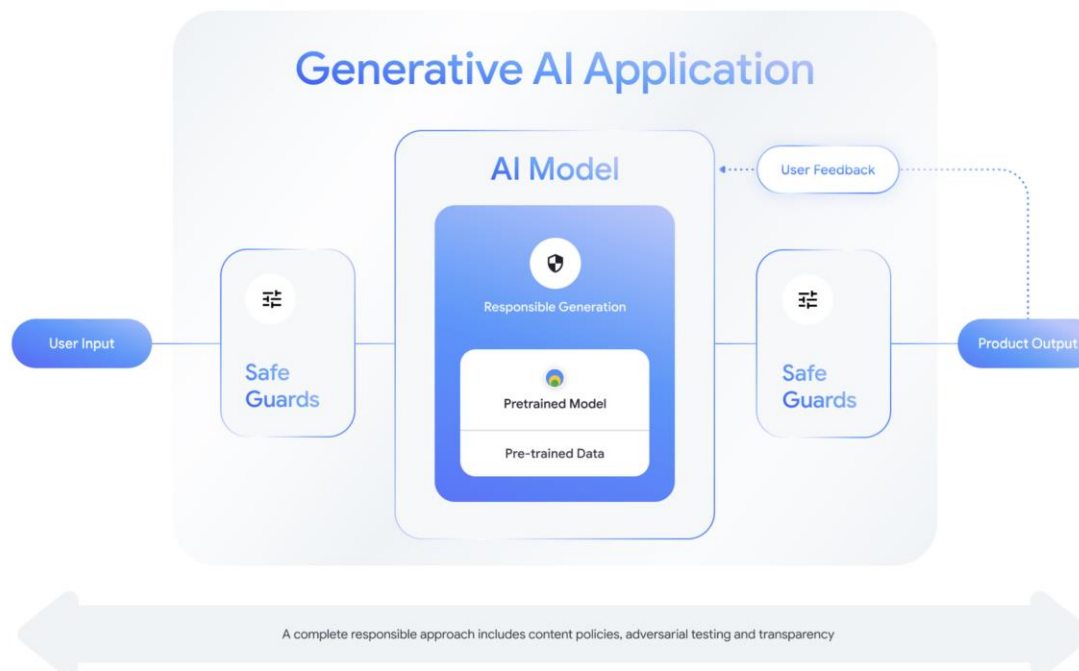
[Send feedback](#)

This toolkit provides resources to apply best practices for responsible use of open models such as the Gemma models, including:

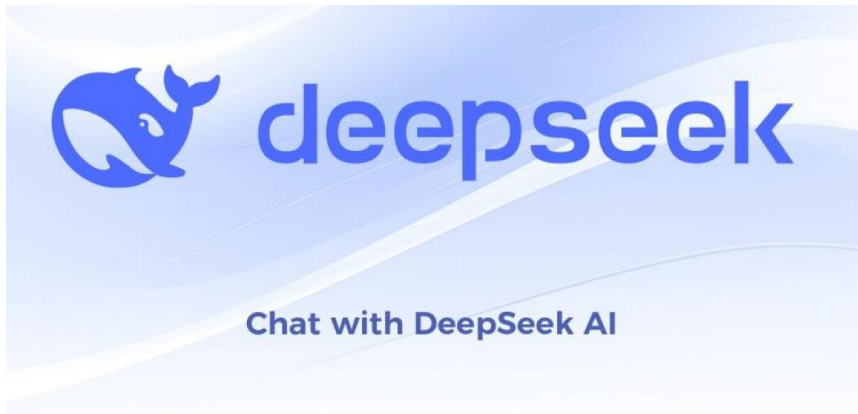
- Guidance on setting safety policies, safety tuning, safety classifiers and model evaluation.
- The [Learning Interpretability Tool \(LIT\)](#) for investigating Gemma's behavior and addressing potential issues.
- A methodology for building robust safety classifiers with minimal examples.

This version of the toolkit focuses on English text-to-text models only. You can provide feedback to make this toolkit more helpful through the feedback mechanism link at the bottom of the page.

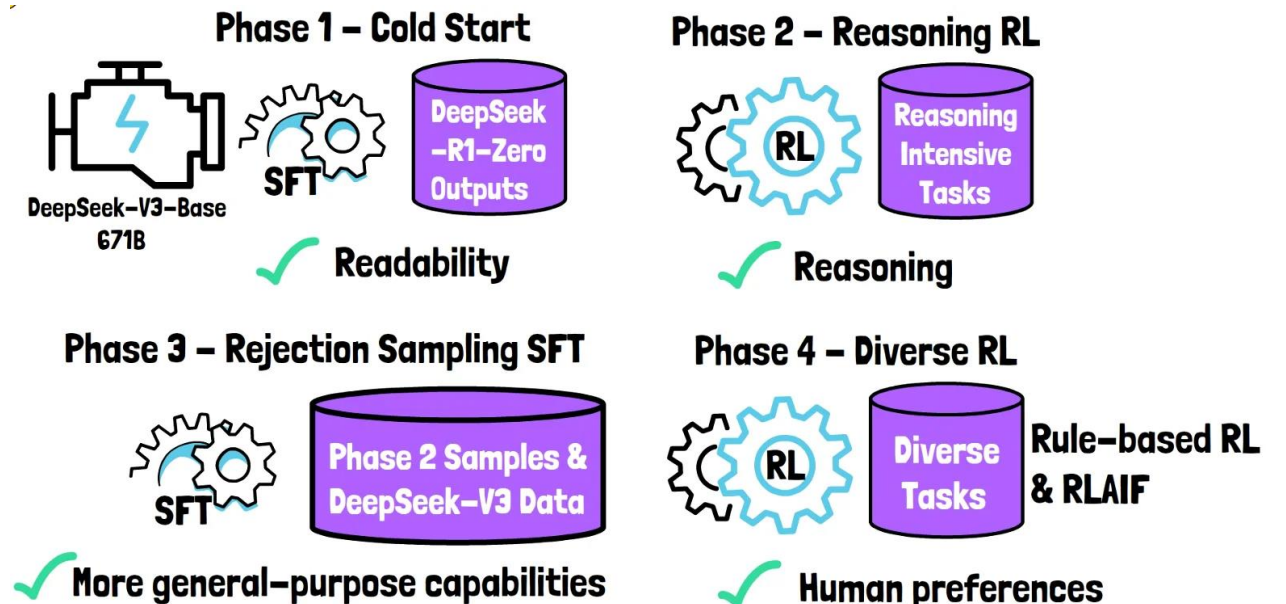
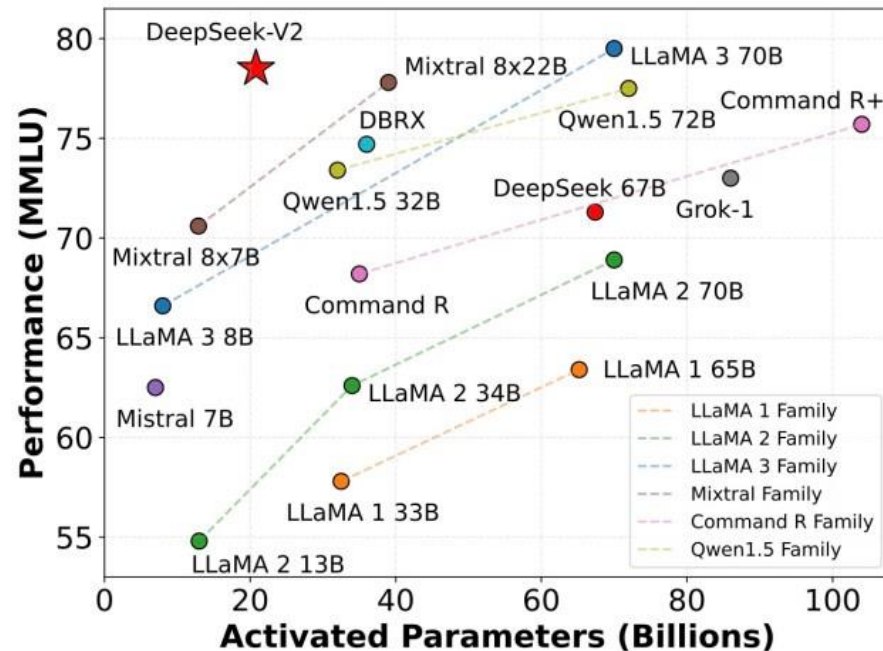
When building with Gemma, you should take a holistic approach to responsibility and consider all the possible challenges at the application and model levels. This toolkit covers risk and mitigation techniques to address safety, privacy, fairness, and accountability.



DeepSeek



DeepSeek is a Chinese AI research lab focused on developing LLMs and AI systems. DeepSeek series includes DeepSeek-V2, a high-performance open-weight model trained on diverse multilingual datasets, aiming to compete with top-tier models like GPT and LLaMA. DeepSeek emphasizes efficiency, scalability, and domain-specific capabilities, particularly in Chinese language understanding and applications. The lab has also developed DeepSeekCoder, an advanced AI model optimized for code generation and programming assistance. With strong computational resources and a growing presence in the AI research landscape, DeepSeek is positioning itself as a significant player in the global LLM ecosystem.



Solutions > Generative AI > Tongyi Qianwen (Qwen)

Tongyi Qianwen (Qwen)

Top-performance foundation models from Alibaba Cloud

[Contact Sales](#)


[Qwen on Model Studio](#)

🔖 Qwen2.5-Max: Exploring the Intelligence of Large-scale MoE Model



 TONGYI
Qwen

 TONGYI
Qwen-VL

 TONGYI
Qwen-Audio

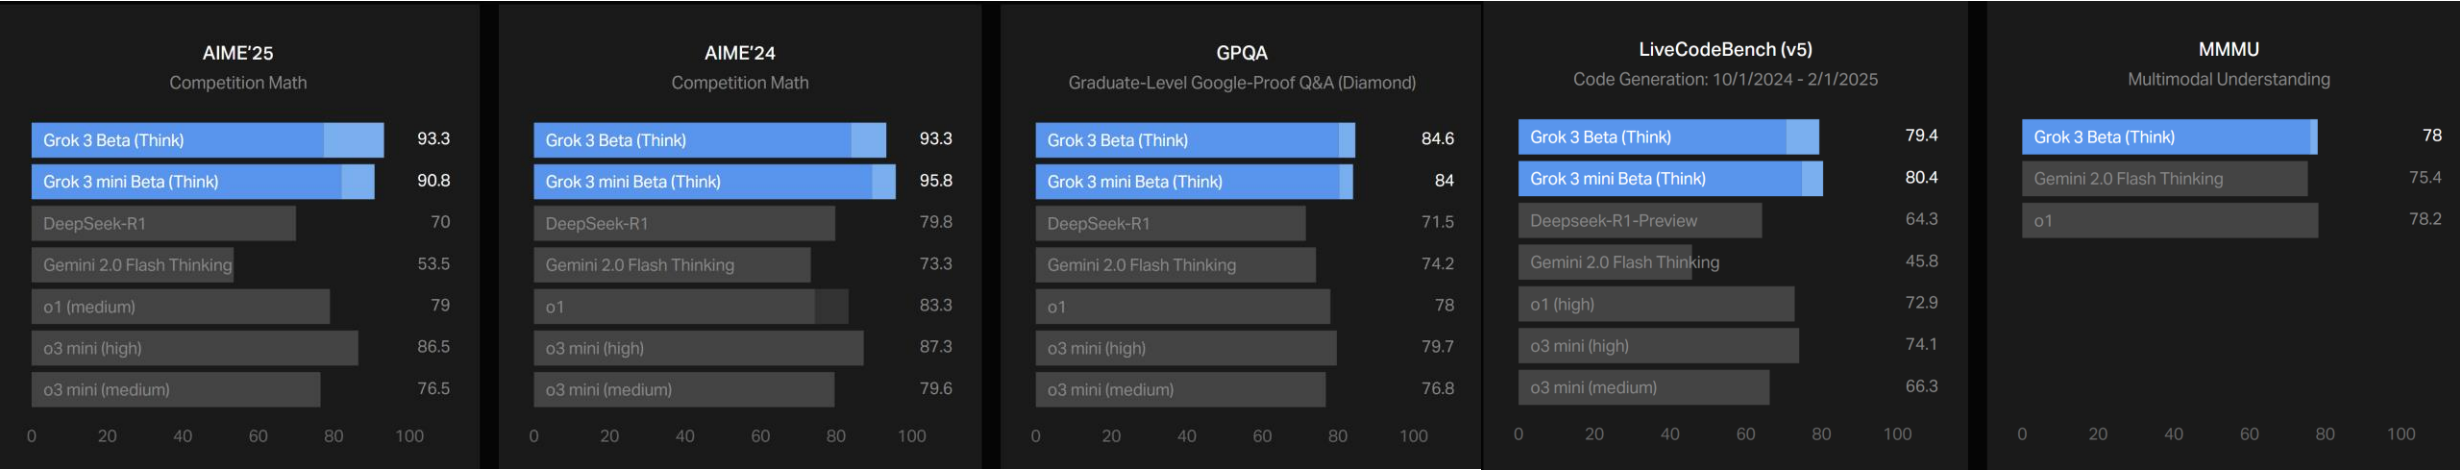
 TONGYI
Qwen2.5-Max

 TONGYI
Qwen 2.5-Coder

 TONGYI
Qwen 2.5-Math

 TONGYI
Qwen-Agent

Grok 3 – Chocolate xAI



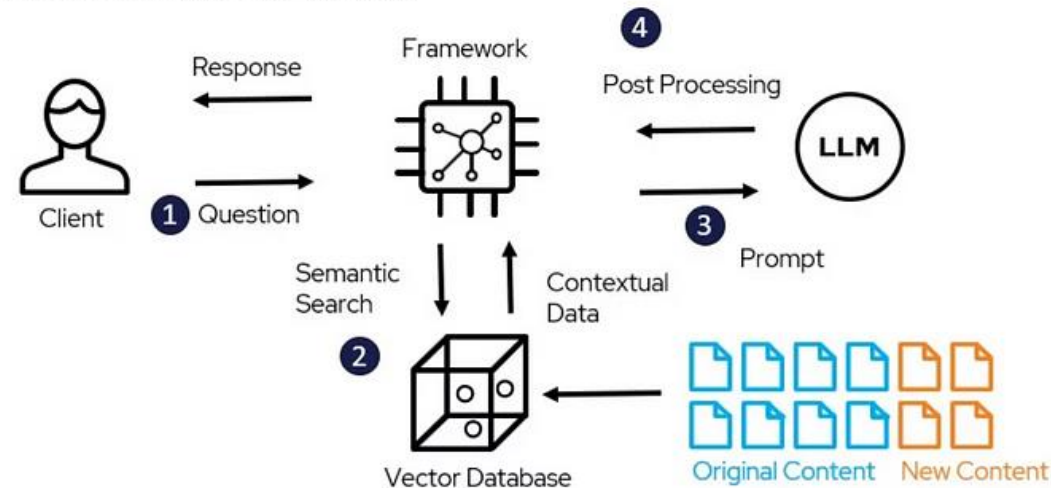
Chatbot Arena LLM Leaderboard (27/02/2025)

| Rank★ (UB) | Rank (StyleCtrl) | Model | Arena Score | 95% CI | Votes | Organization | License |
|------------|------------------|---|-------------|--------|-------|--------------|-------------|
| 1 | 1 | chocolate (Early Grok-3) | 1403 | +6/-6 | 9992 | xAI | Proprietary |
| 2 | 3 | Gemini-2.0-Flash-Thinking-Exp-01-21 | 1385 | +4/-6 | 15083 | Google | Proprietary |
| 2 | 3 | Gemini-2.0-Pro-Exp-02-05 | 1380 | +5/-6 | 13000 | Google | Proprietary |
| 2 | 1 | ChatGPT-4o-latest (2025-01-29) | 1377 | +5/-5 | 13470 | OpenAI | Proprietary |
| 5 | 3 | DeepSeek-R1 | 1362 | +7/-7 | 6581 | DeepSeek | MIT |
| 5 | 8 | Gemini-2.0-Flash-001 | 1358 | +7/-7 | 10862 | Google | Proprietary |
| 5 | 3 | o1-2024-12-17 | 1352 | +5/-5 | 17248 | OpenAI | Proprietary |
| 8 | 7 | o1-preview | 1335 | +3/-4 | 33169 | OpenAI | Proprietary |
| 8 | 8 | Qwen2.5-Max | 1334 | +5/-5 | 9282 | Alibaba | Proprietary |
| 8 | 7 | o3-mini-high | 1332 | +5/-9 | 5954 | OpenAI | Proprietary |
| 11 | 11 | DeepSeek-V3 | 1318 | +4/-5 | 19461 | DeepSeek | DeepSeek |
| 11 | 13 | Qwen-Plus-0125 | 1311 | +9/-7 | 5112 | Alibaba | Proprietary |
| 11 | 14 | GLM-4-Plus-0111 | 1310 | +6/-9 | 5134 | Zhipu | Proprietary |

RAG Architecture Model

- Retrieval Augmented Generation (RAG) has emerged as a powerful technique for improving LLMs
- Retrieving and conditioning external knowledge, RAG allows models to generate more accurate, relevant, and comprehensive text

RAG Architecture Model



<https://medium.com/@gulernilay088/>

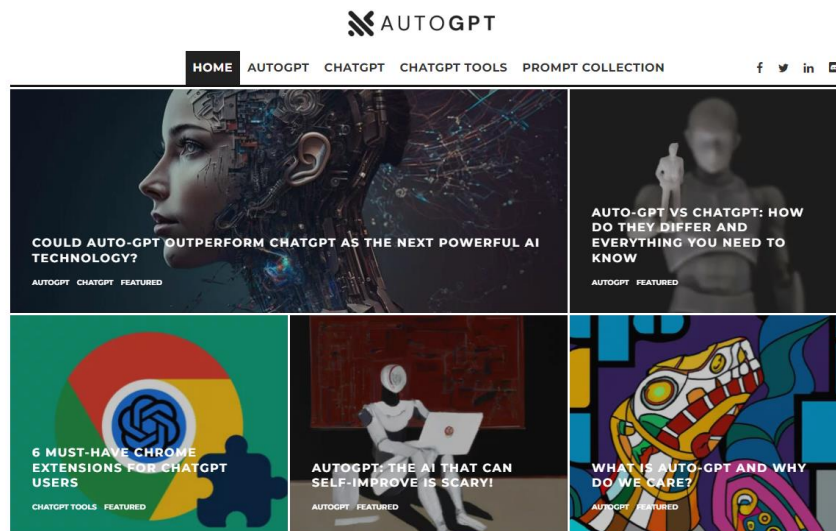
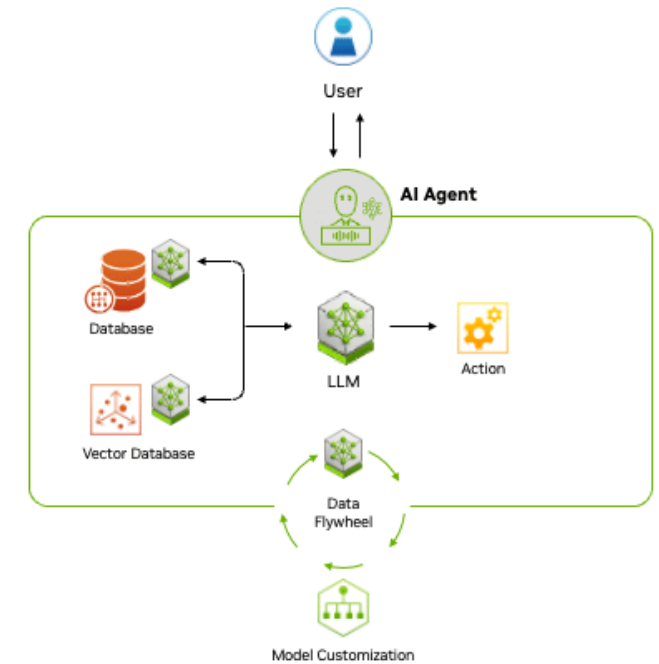
- Advanced RAG enhances each module further with innovations like higher-order retrievers, cross-encoder rerankers, and evidence manipulation architectures

AI & LLM Key Issues in 2024/25

- **Small Language Models (SLM)** will drive **enterprise AI adoption**
- Mixture of Experts (**MoE**) + **LoRA** will enable SLMs to outperform 10x larger LLMs
- **Open-source will become the de facto** way for **LLMs** use
- Large language models will adopt **modular architectures**
- **Synthetic data** revolution arrives to LLMs
- LLM **hallucinations disappear** as training techniques evolve
- **Data** (not LLMs) becomes the **true competitive advantage**
- LLM adoption grows as the Transformer architecture matures
- **RAGs** (Retrieval-Augmented Generation) used in most LLMs tasks
- Shift from closed to open-source **democratizes AI** (transparency)
- **LLMs for Machines and Robots**, not just people

Auto-GPT – AI Agents

- Auto-GPT is an "AI agent" that given a goal in natural language, will attempt to achieve it by breaking it into sub-tasks and using the internet and other tools in an automatic way using LLMs APIs to perform autonomous tasks
- The difference between an agent and a language model is that agents complete task autonomously



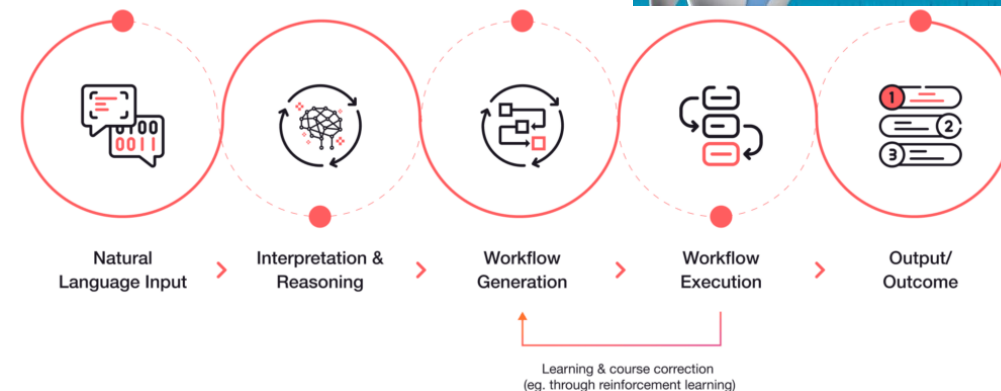
Features

- Internet access for searches and information gathering
- Long-term and short-term memory management
- GPT-4 instances for text generation
- Access to popular websites and platforms
- File storage and summarization with GPT-3.5
- Extensibility with Plugins

Agentic AI

- **Autonomy:** Ability to initiate and complete tasks with limited direct human supervision, with great flexibility and efficiency in executing tasks
- **Reasoning:** Sophisticated decision-making based on context and trade-offs
- **Reinforced Learning:** Dynamically evolve by interacting with their environment and receiving feedback from these interactions.
- **Language Understanding:** Capacity to understand and follow complex instructions
- **Workflow Optimization:** Efficient execution of multi-step complex processes

| Feature | Agentic AI | Generative AI | Traditional AI |
|------------------|---|--|---|
| Primary Function | Goal-oriented action & decision-making | Content generation (text, code, images, etc.) | Focused on automating repetitive tasks |
| Autonomy | High – Operates with minimal human oversight | Variable – May require user prompts or guidance | Low – Relies on specific algorithms and set rules |
| Learning | Reinforced Learning – Improves through experience | Data-driven learning – Learns from existing data | Relies on predefined rules and human intervention |



<https://aisera.com/blog/agentic-ai/>

Agent-Based Simulation

- **Simulation:** Imitation of some real thing, state of affairs, or process, over time, representing certain **key characteristics** or behaviours of the physical or abstract system
- Applications:
 - Understand system **functioning**
 - **Performance optimization**
 - Testing and validation
 - **Decision making**
 - Training and education
 - **Test future/expensive systems**
- Complex systems impossible to solve mathematically
- Agent Based Modeling and Simulation
- **Compress/Accelerate Time: Machine Learning**

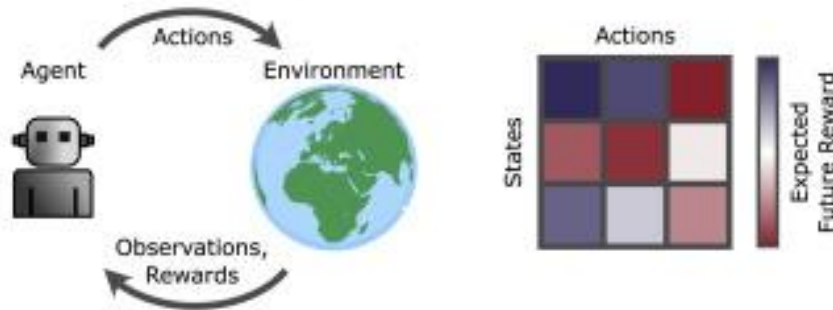


Deep Reinforcement Learning (DRL)

A Classic Reinforcement Learning

Reinforcement Learning Problem

Tabular Solution



B Classic Deep Learning

Categorization Problem

Deep Learning Solution

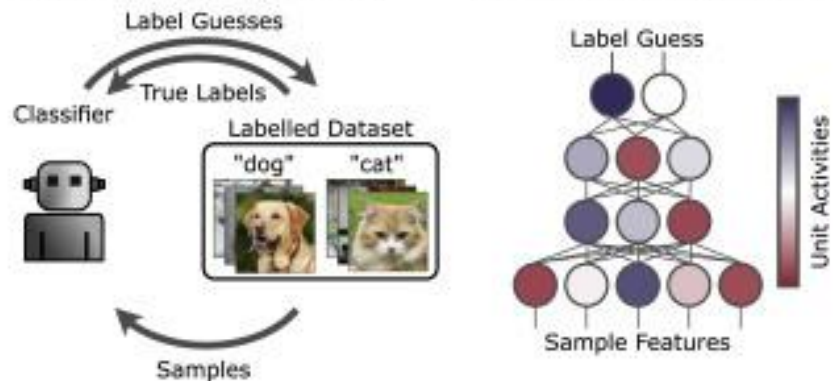
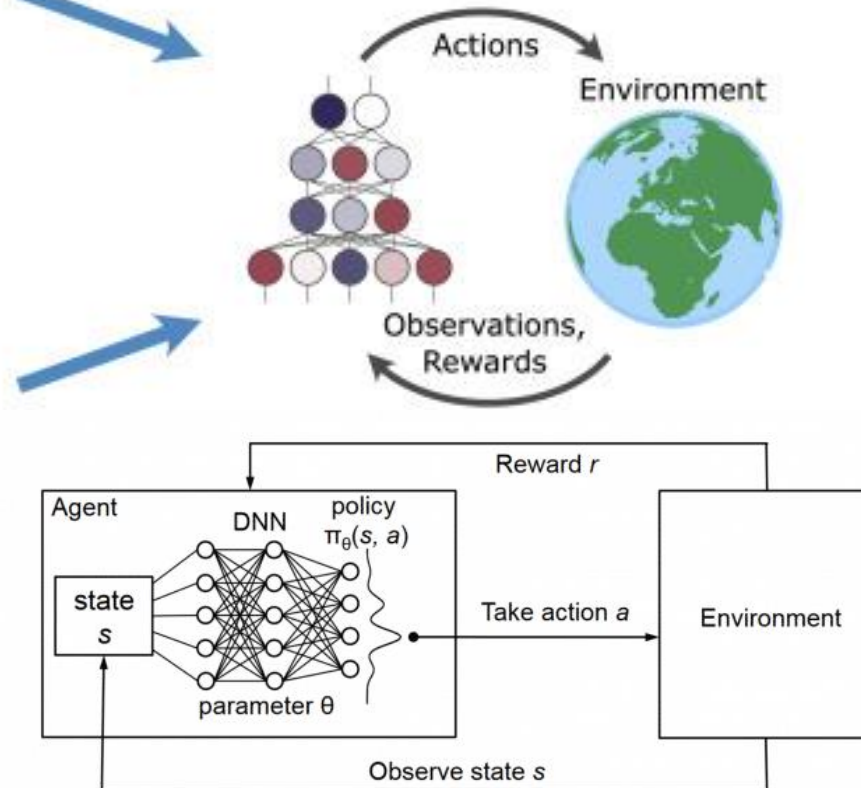


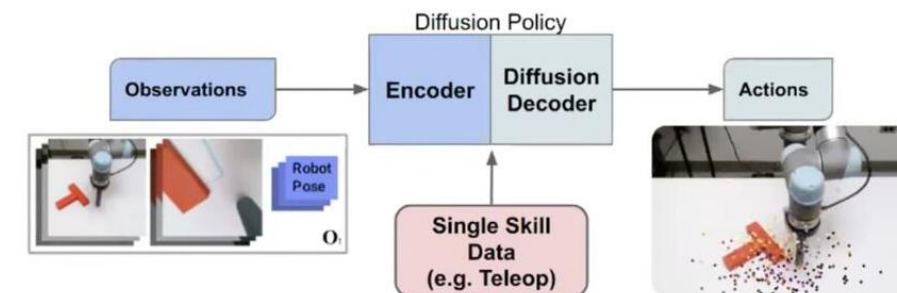
Image: <https://ucsdneuro.wordpress.com/2021/06/13/diving-deep-into-the-brain-a-case-for-deep-reinforcement-learning-in-neuroscience/>

C Deep Reinforcement Learning: Deep learning solutions for RL problems



Large Behavior Models - LBMs

- **Large Behaviour Models** are AI systems that go beyond processing language. They're designed to **interpret, predict, and generate complex sequences of human-like actions and behaviors**
- While LLMs focus on text and dialogue, LBMs specialize in **behavioural patterns, decision-making processes, and interactive scenarios**
- LLMs try to guess the next word, while **LBMs try to guess the next action or behaviour**
- LBMs are enhancing robots' abilities to move naturally and respond intuitively to their surrounding
- **LBMs** rely on:
 - **Transformer Architecture:** Similar to LLMs but adapted to handle sequences of actions/behaviours
 - **Multi-Modal Learning:** Integrating visual, spatial, and temporal data to enhance understanding
 - **Reinforcement Learning:** Learning from interactions and feedback to improve behavior over time



Transformers LBM's and DRL: RT-X



https://www.youtube.com/watch?v=a7BkJmTbohl&ab_channel=AIRevolution

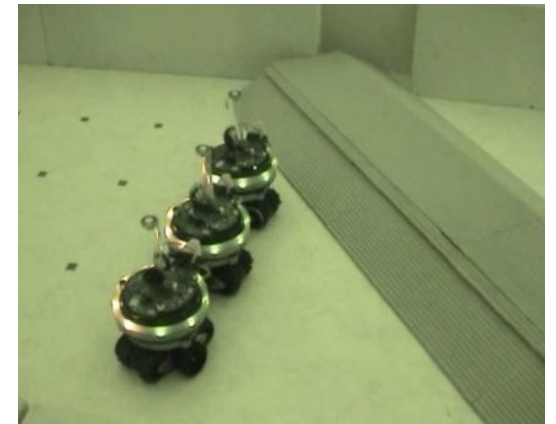
Coordination in Multi-Agent Systems

Motivation:

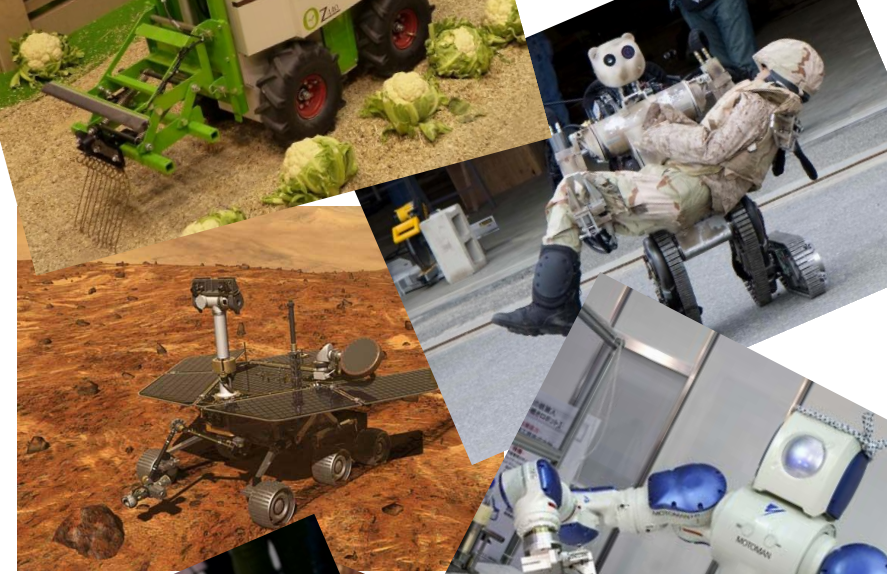
- Agents don't live alone and have to work in a group...
- **Human-Computer Interaction**
- **Multi-Agent Coordination**

Coordination : “to work in harmony in a group”

- **Dependencies** in agent actions
- Global constraints
- **No agent**, individually **has enough resources**, information or capacity to execute the task or solve the problem
- **Efficiency**: Information exchange or tasks division
- **Prevent anarchy and chaos**: Partial vision, lack of authority, conflicts, agent's interactions



AI in Robotics



Robotic Games and Competitions

Benefits

- Research inspiration
- Hard deadline for creating fully functional system
- Common platform/problem
- Exchange of research ideas/solutions
- Continually improving solutions
- Excitement for students/researchers at all I
- Large number of teams/solutions created
- Encouragement for flexible software/hardware

Dangers

- Obsession with winning
- Domain dependent/ hacked solutions
- Cost escalation
- Difficulty in entering at competitive level
- Restrictive rules
- Invalid evaluation conclusions

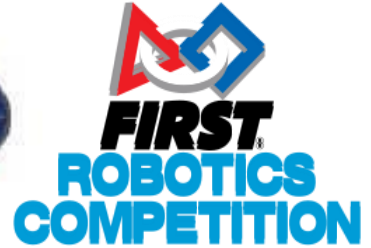


ROBOXOTICA
Festival für Cocktail-Robotik



LAWRENCE TECHNOLOGICAL UNIVERSITY
ROBOFEST

ROBOTICA
FESTIVAL NACIONAL



EUROPEAN
ROBOTICS
LEAGUE



RoboCup: Objectives

- Joint International Project:
 - (Distributed) Artificial Intelligence
 - Intelligent Robotics
- Soccer – Central Research Topic:
 - Very complex collective game
 - Huge number of technologies involved:
 - Autonomous Agents, Multi-Robot Systems, Cooperation, Communication, Strategic Reasoning, Robotics, Sensor Fusion, Real-Time Reasoning, Machine Learning, etc
- 5 Challenges: Soccer, Rescue, @Home, Industrial, Junior
- 17 Leagues, 2500 Humans and 2000 Robots!
- Main Goal: ***“By 2050, develop a team of fully autonomous humanoid robots that may win against the human world champion team in soccer!”***



RoboCup Leagues and Challenges

Industrial



Mobility (Remote Operator)



Stairs with Debris



Rescue

@Home

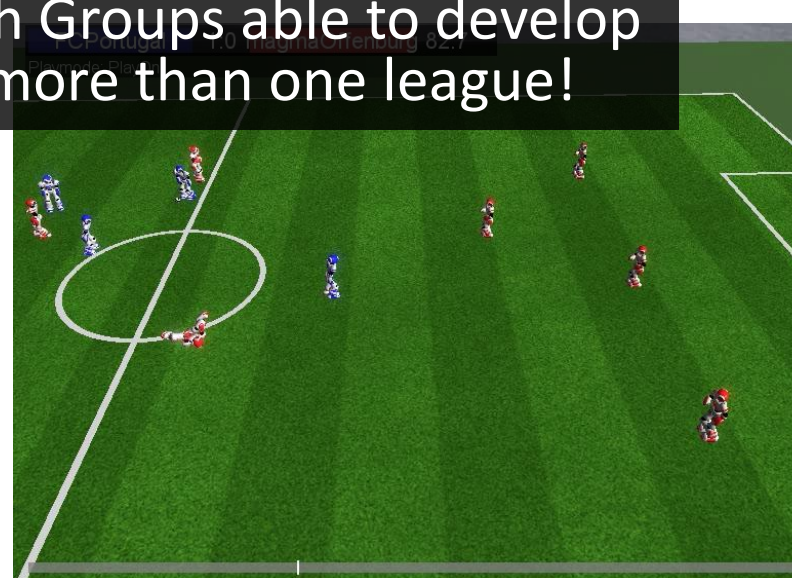


Soccer

Robotic Soccer Competitions

RoboCup

- Real, Standard, Simulated Robots
- Mini, Small, Medium and Large Robots
- Wheeled, Legged and Humanoid Robots
- **Distinct but interrelated Leagues/Problems**
- Only a Few Research Groups able to develop code that works in more than one league!



Robotic Soccer Competitions

Middle
Size
League



Standard
Platform
League



Small
Size
League

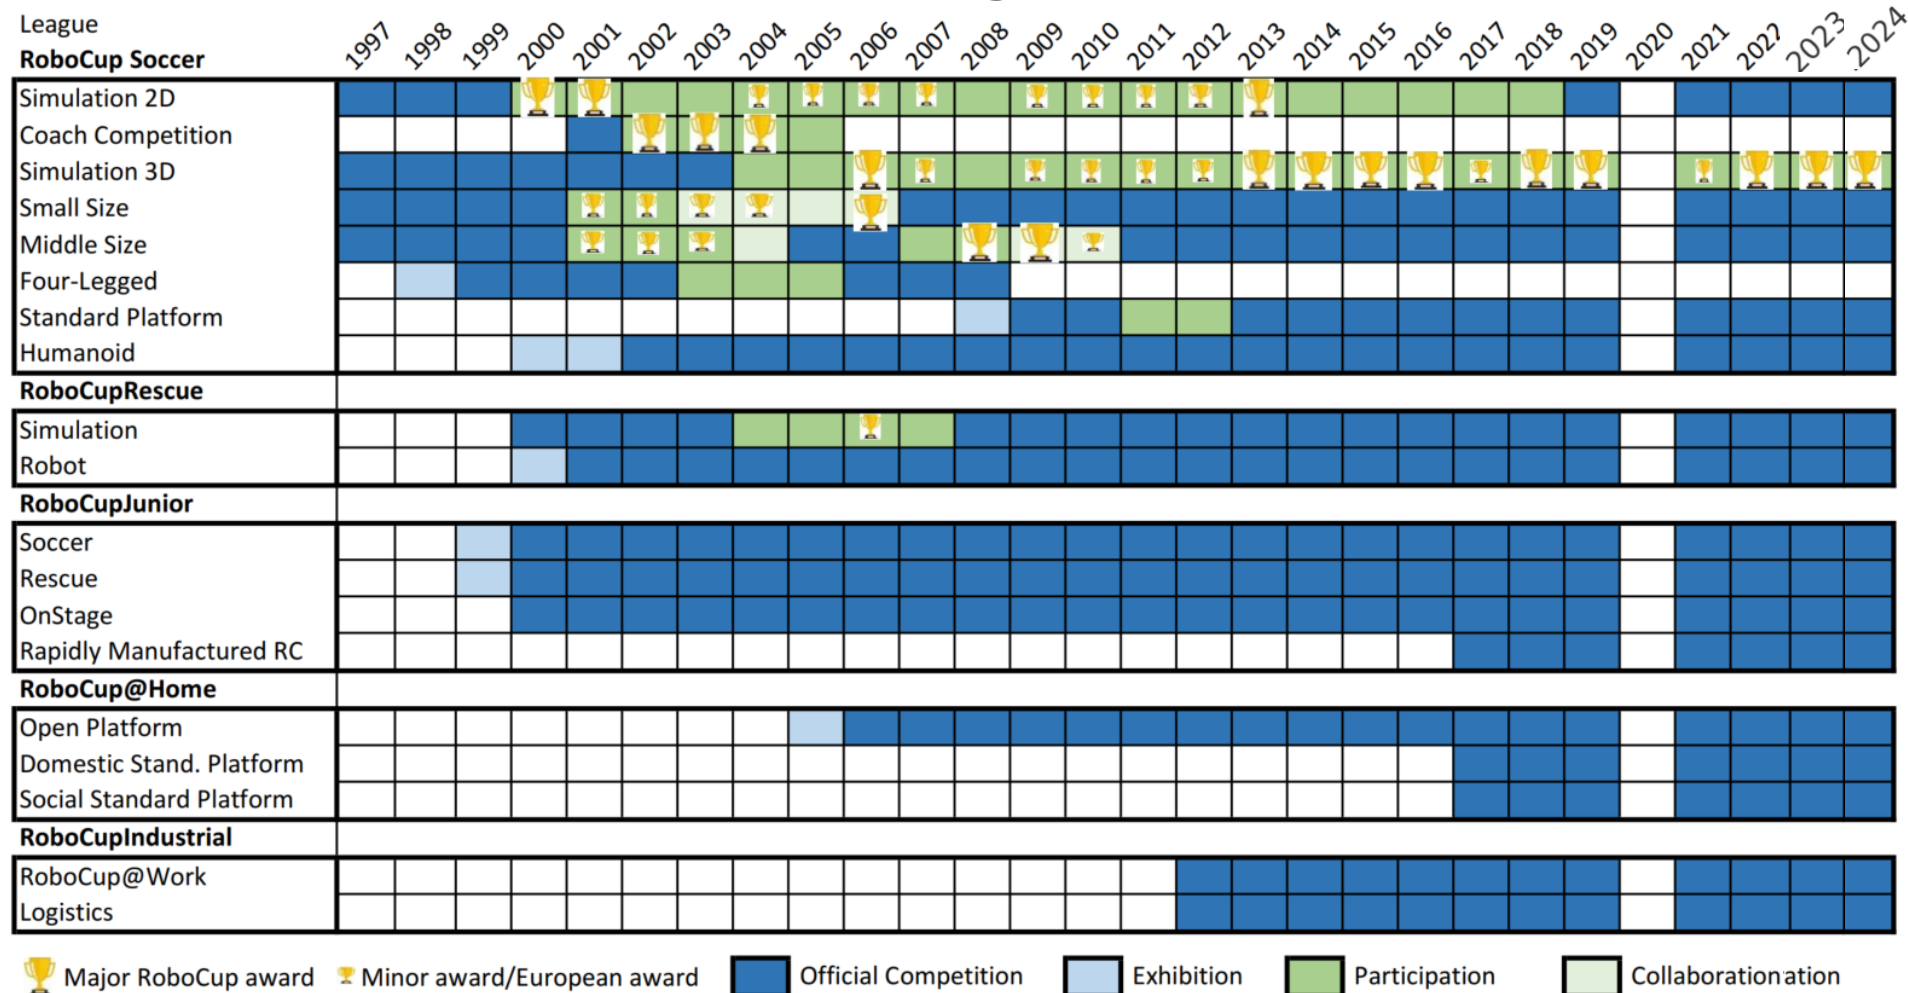


Humanoid
League



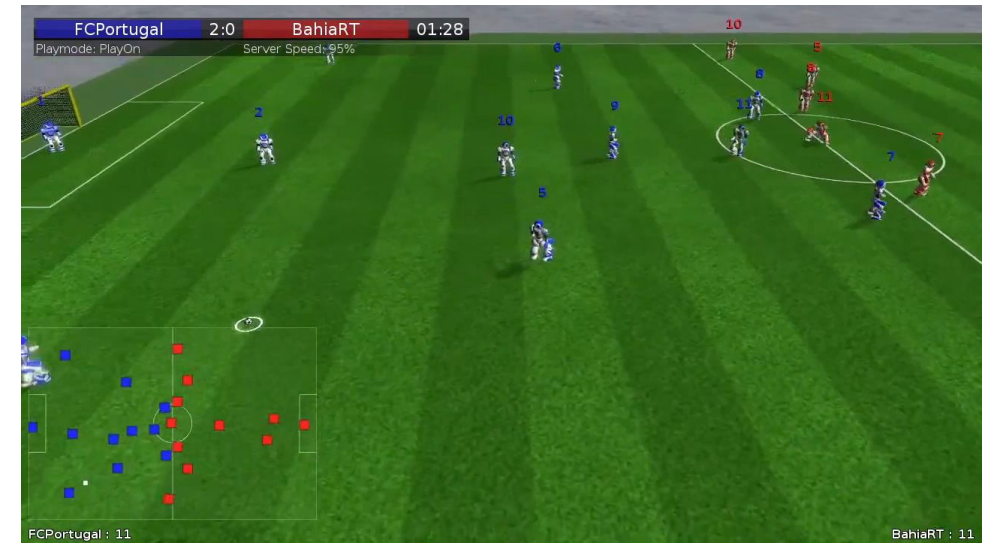
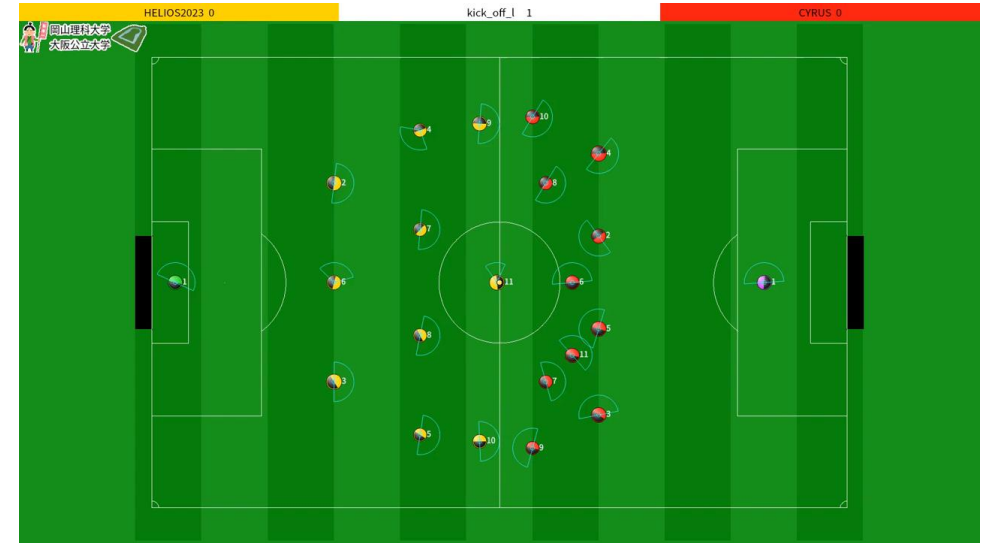
RoboCup Leagues and Participations

FC Portugal Team



Simulation 3D League (Humanoids)

- **Realistic physics and humanoid robot model:**
 - Spheres in 2004, Humanoids in 2007, NAO Robot Model: 2008, Heterogeneous Robots: 2013
- **Strong relation with SPL and Humanoid**
- 2 vs 2 -> 6 vs 6 -> 9 vs 9 -> 11 vs 11
- **Server/Simulator (SimSpark)**
 - Updates world state
 - Forces the “**laws of physics**”: collisions, drag, gravity, ...
 - Send sensor information (**perceptors**)
 - Executes actions (**effectors**)
 - Enforces **soccer rules** – referee
- **Impossible to create competitive skills by hand!**



Simulation 3D League (Humanoids)



Simulation 3D League Challenges

- **Real-time, realistic, multi-robot simulation**
- Several sources of sensorial information: **visual, auditory, and physical**
- **Unreliable** and low-bandwidth **communication**
- **Heterogeneous robots** (5 different types)
- Impossible to generate **competitive skills** by hand
- **Multi-objective**, partially **cooperative**, partially **competitive environment**
- Need to **transform** very **low-level actions** (use motors to move joints in time) **into high-level skills** (walking, kicking, dribbling)
- Need to create **complex collective actions** (passes, setplays, formations)
- **Evolving rules** and challenges

Deep Reinforcement Learning (DRL)

Model-Free RL learns directly from interaction with the environment.

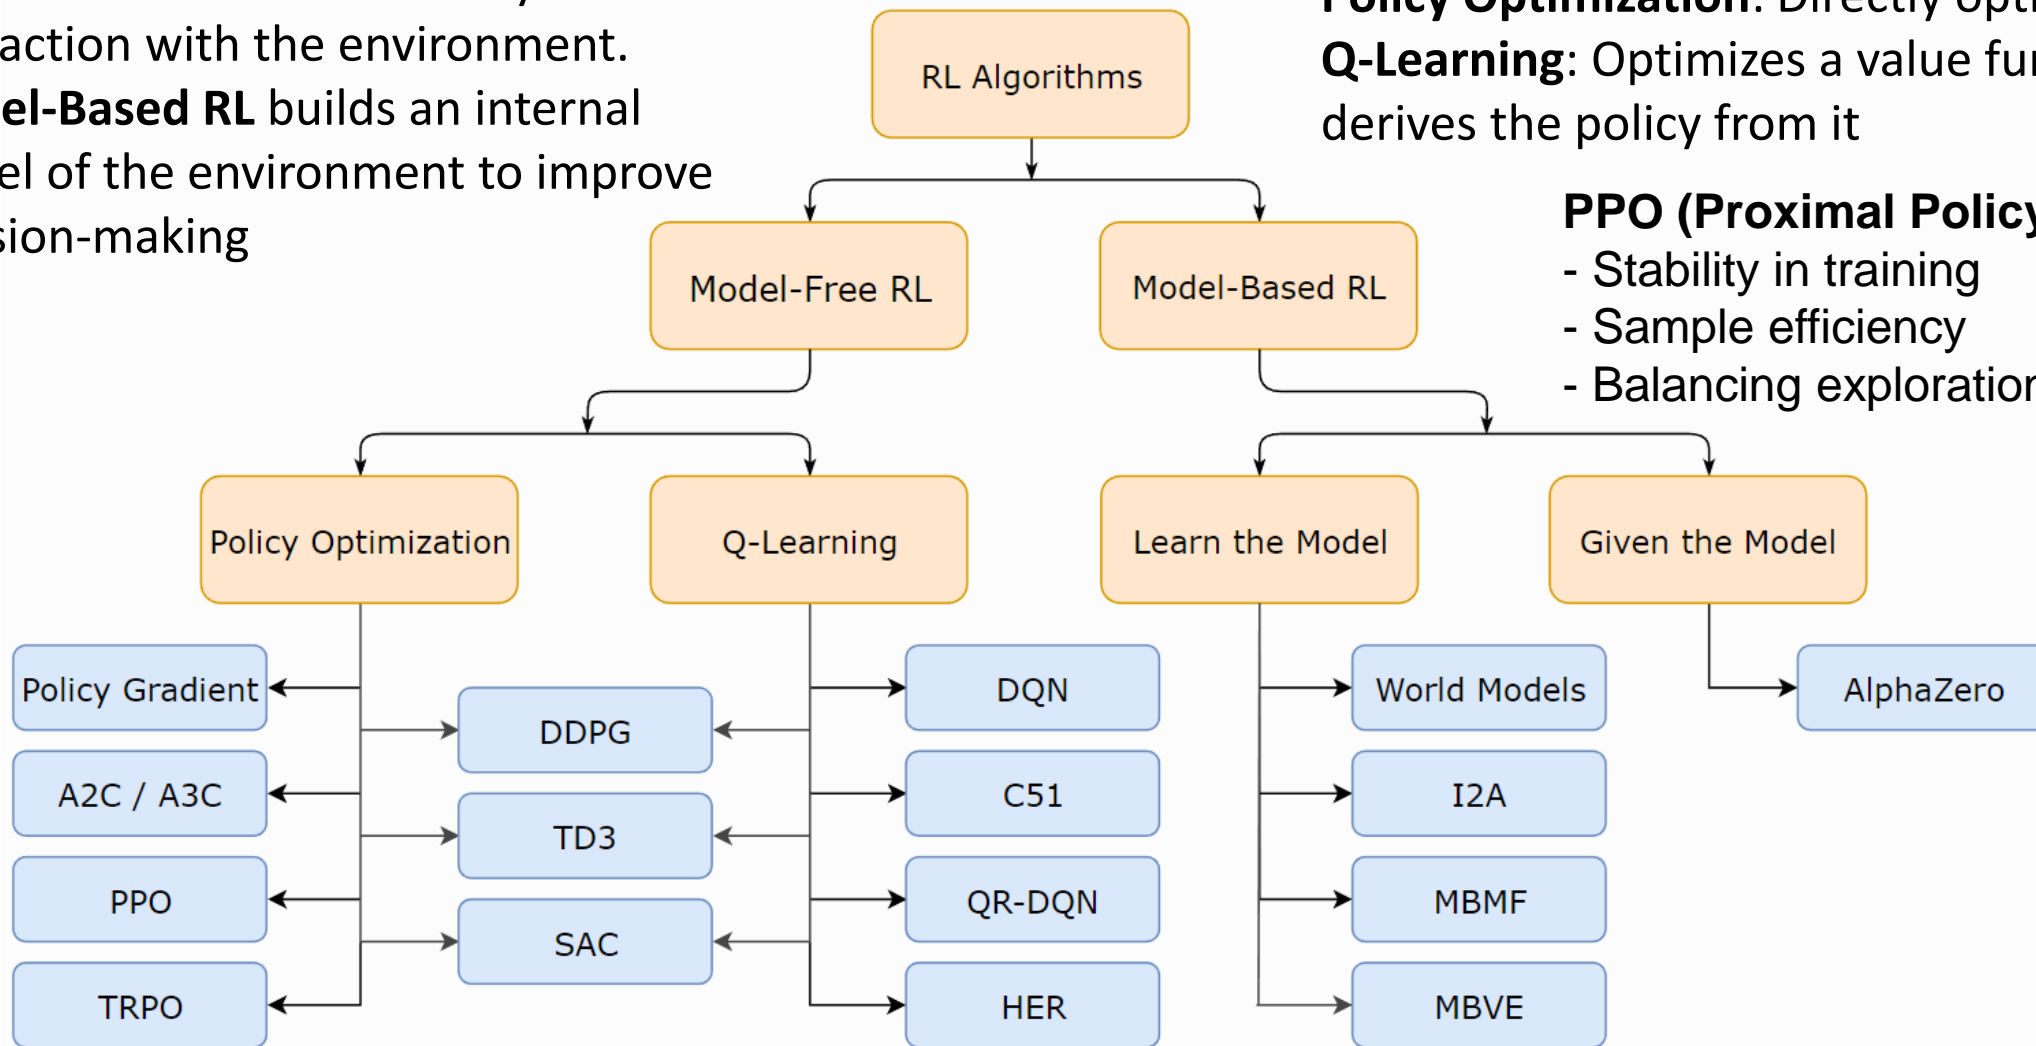
Model-Based RL builds an internal model of the environment to improve decision-making

Policy Optimization: Directly optimizes policy

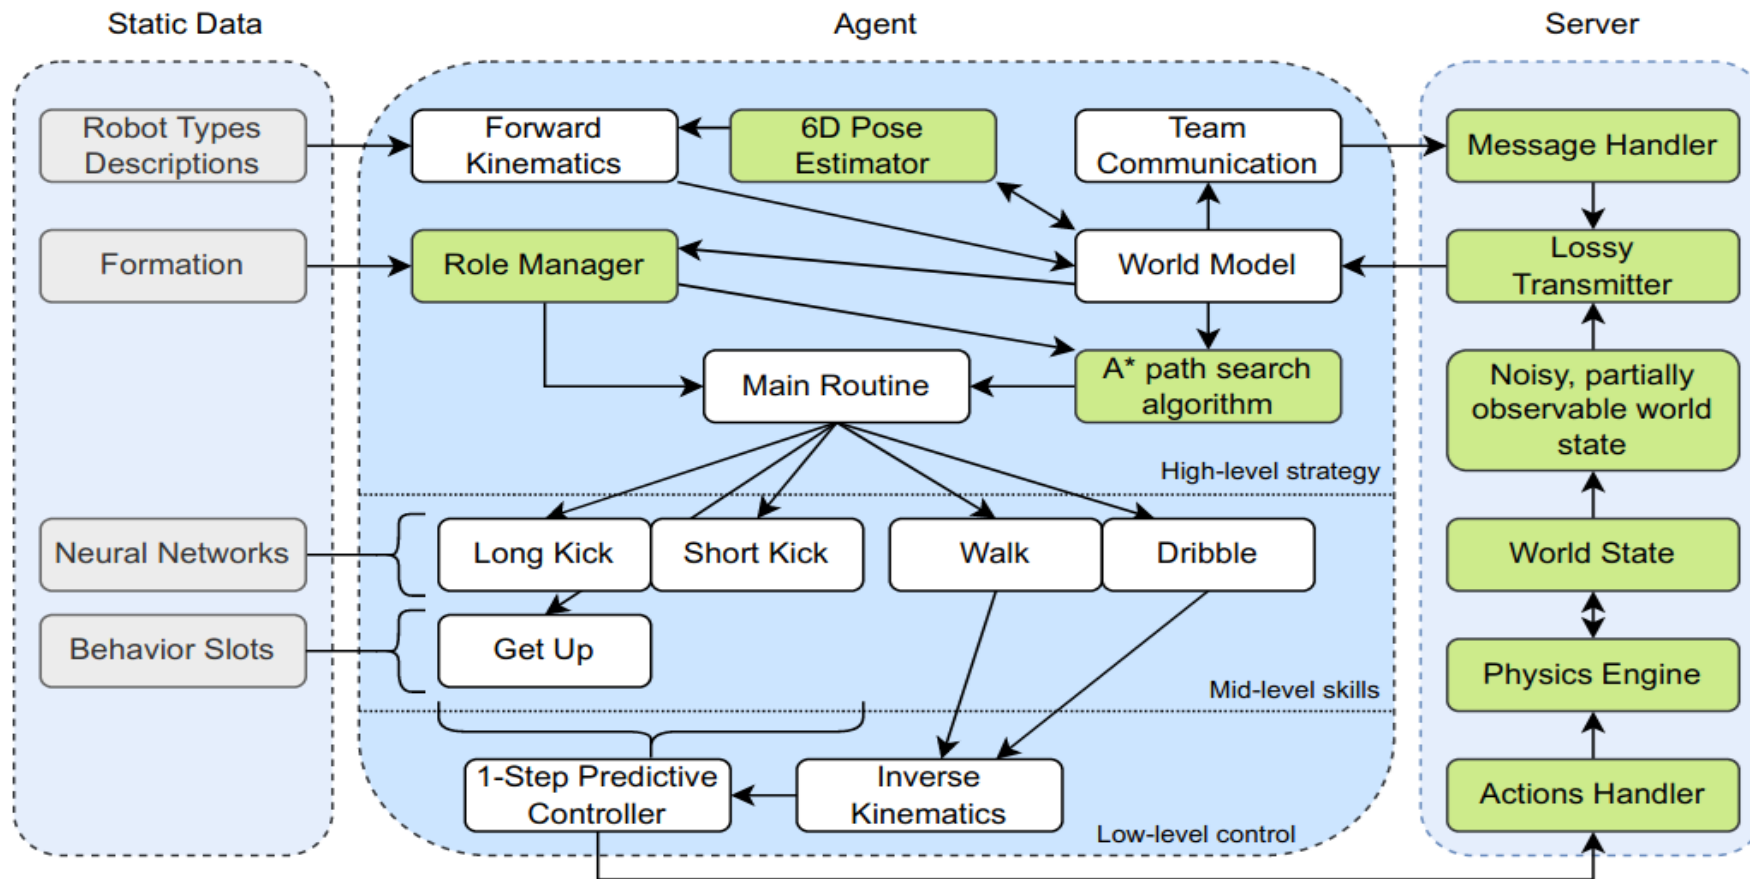
Q-Learning: Optimizes a value function and derives the policy from it

PPO (Proximal Policy Optimization)

- Stability in training
- Sample efficiency
- Balancing exploration and exploitation



Robot/Team Architecture

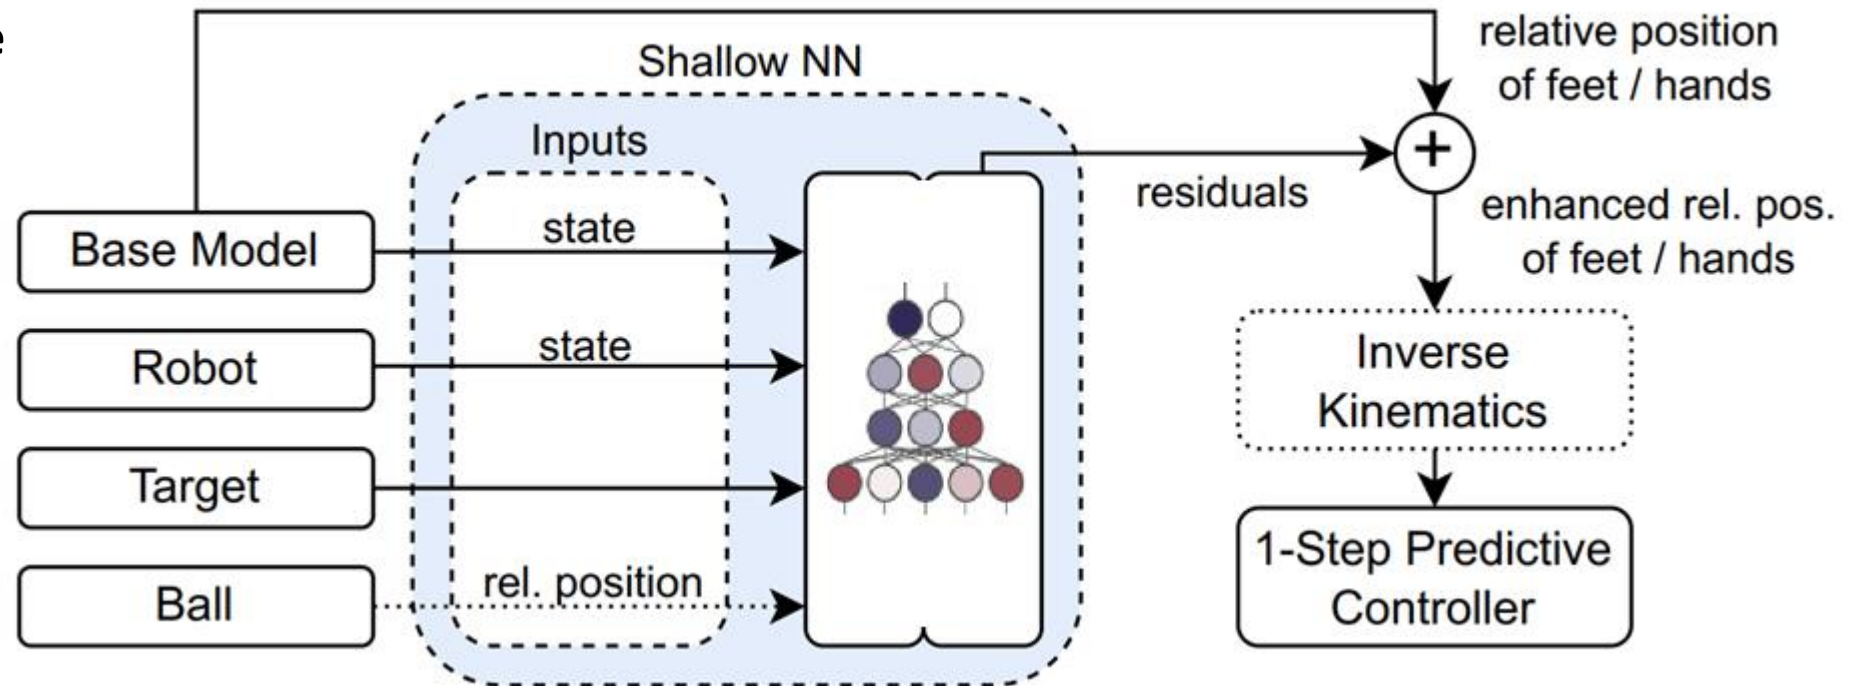


- Short kick (3 - 9m) for passes and long kick (17 - 19m) for shooting
- Omnidirectional walk (0.7 - 0.9 m/s)
- Dribble (1.2 - 1.4 m/s)
- Get Up (3 variations per robot type: front, back, side)

Model Architecture

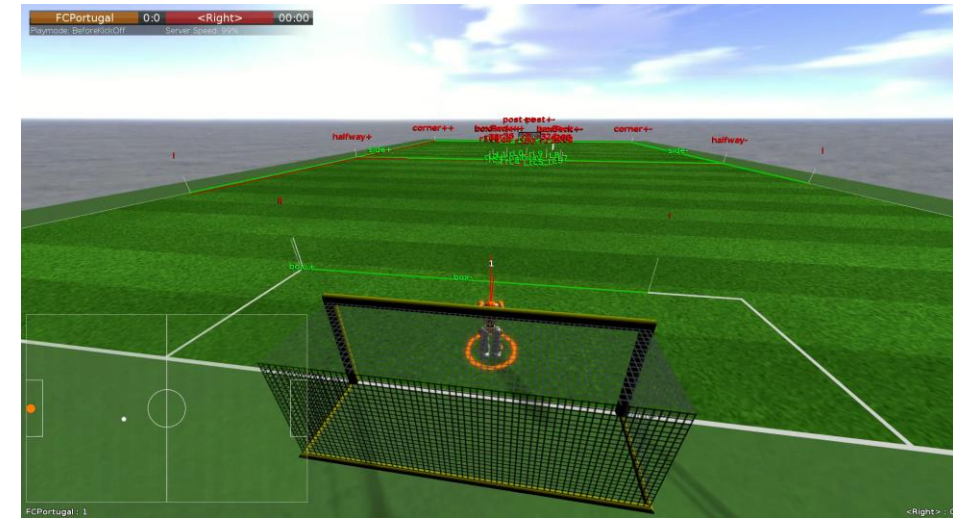
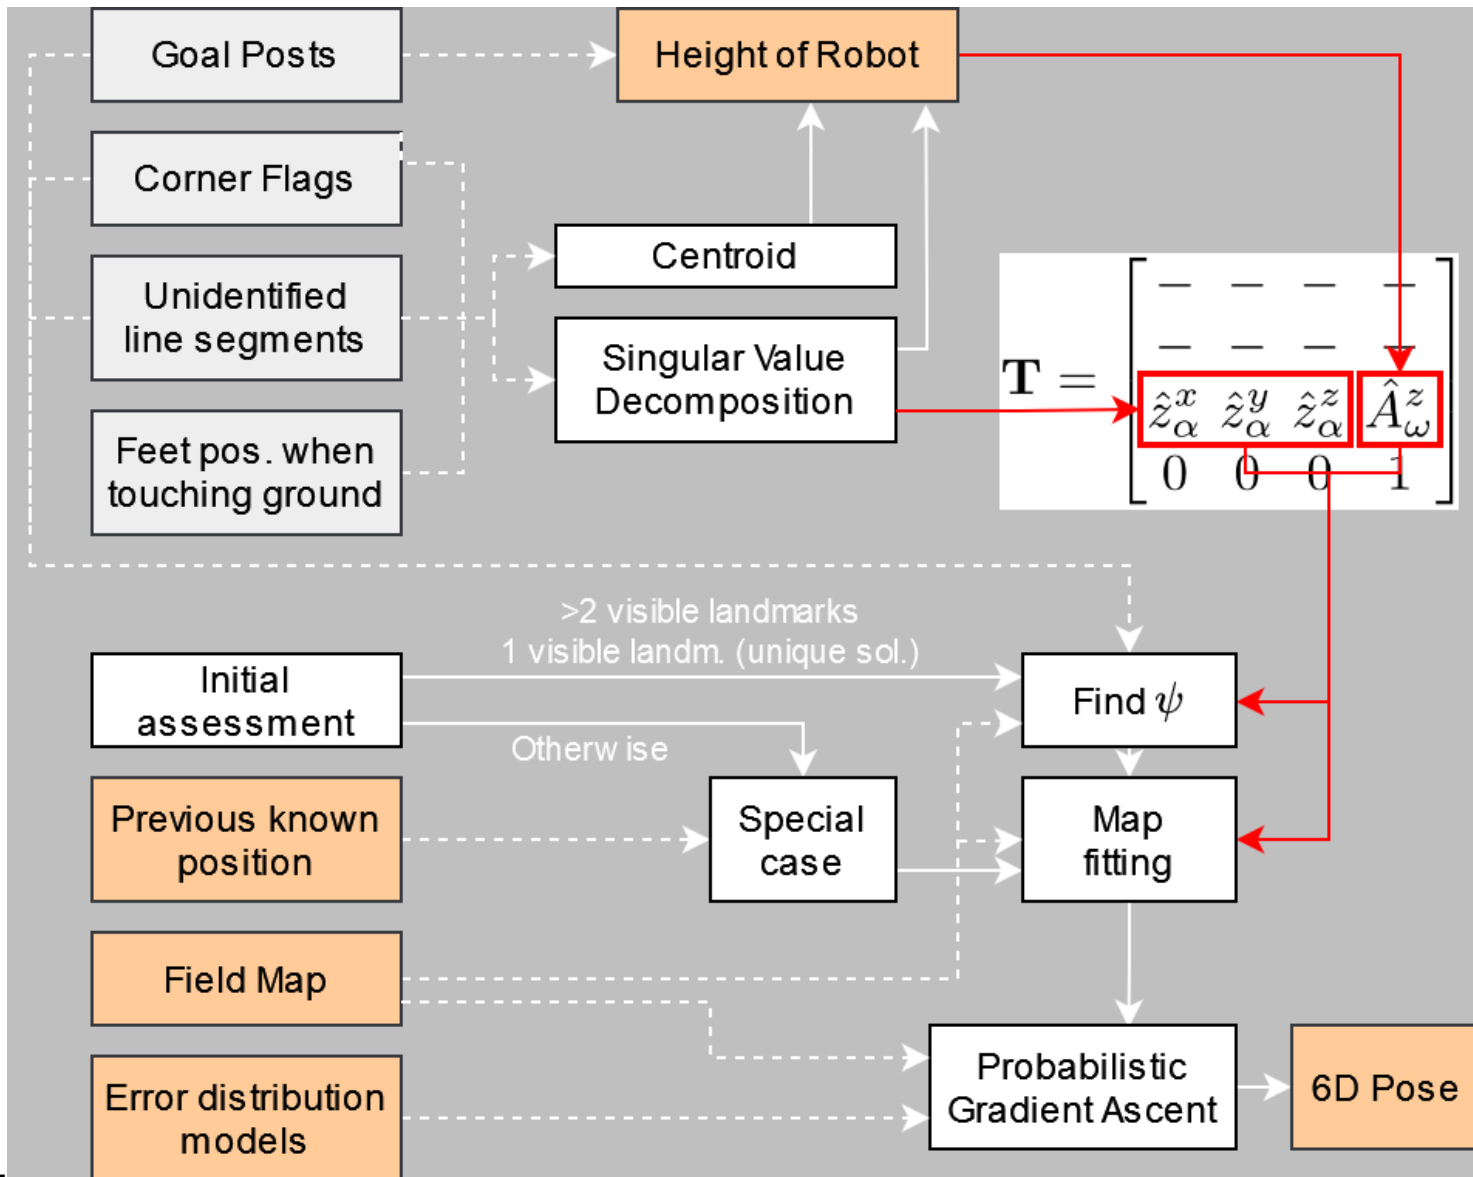
- Enhanced motion control pipeline for a robot, integrating **neural networks** and **predictive control**
- Learning-based and physics-based approaches.
- NN acts as an **enhancement layer**, improving traditional motion models
- **Precise movements**, such as **kicking** or **dribbling** a ball

Model Architecture for the Long/Short Kick, Walk, and Dribble skills



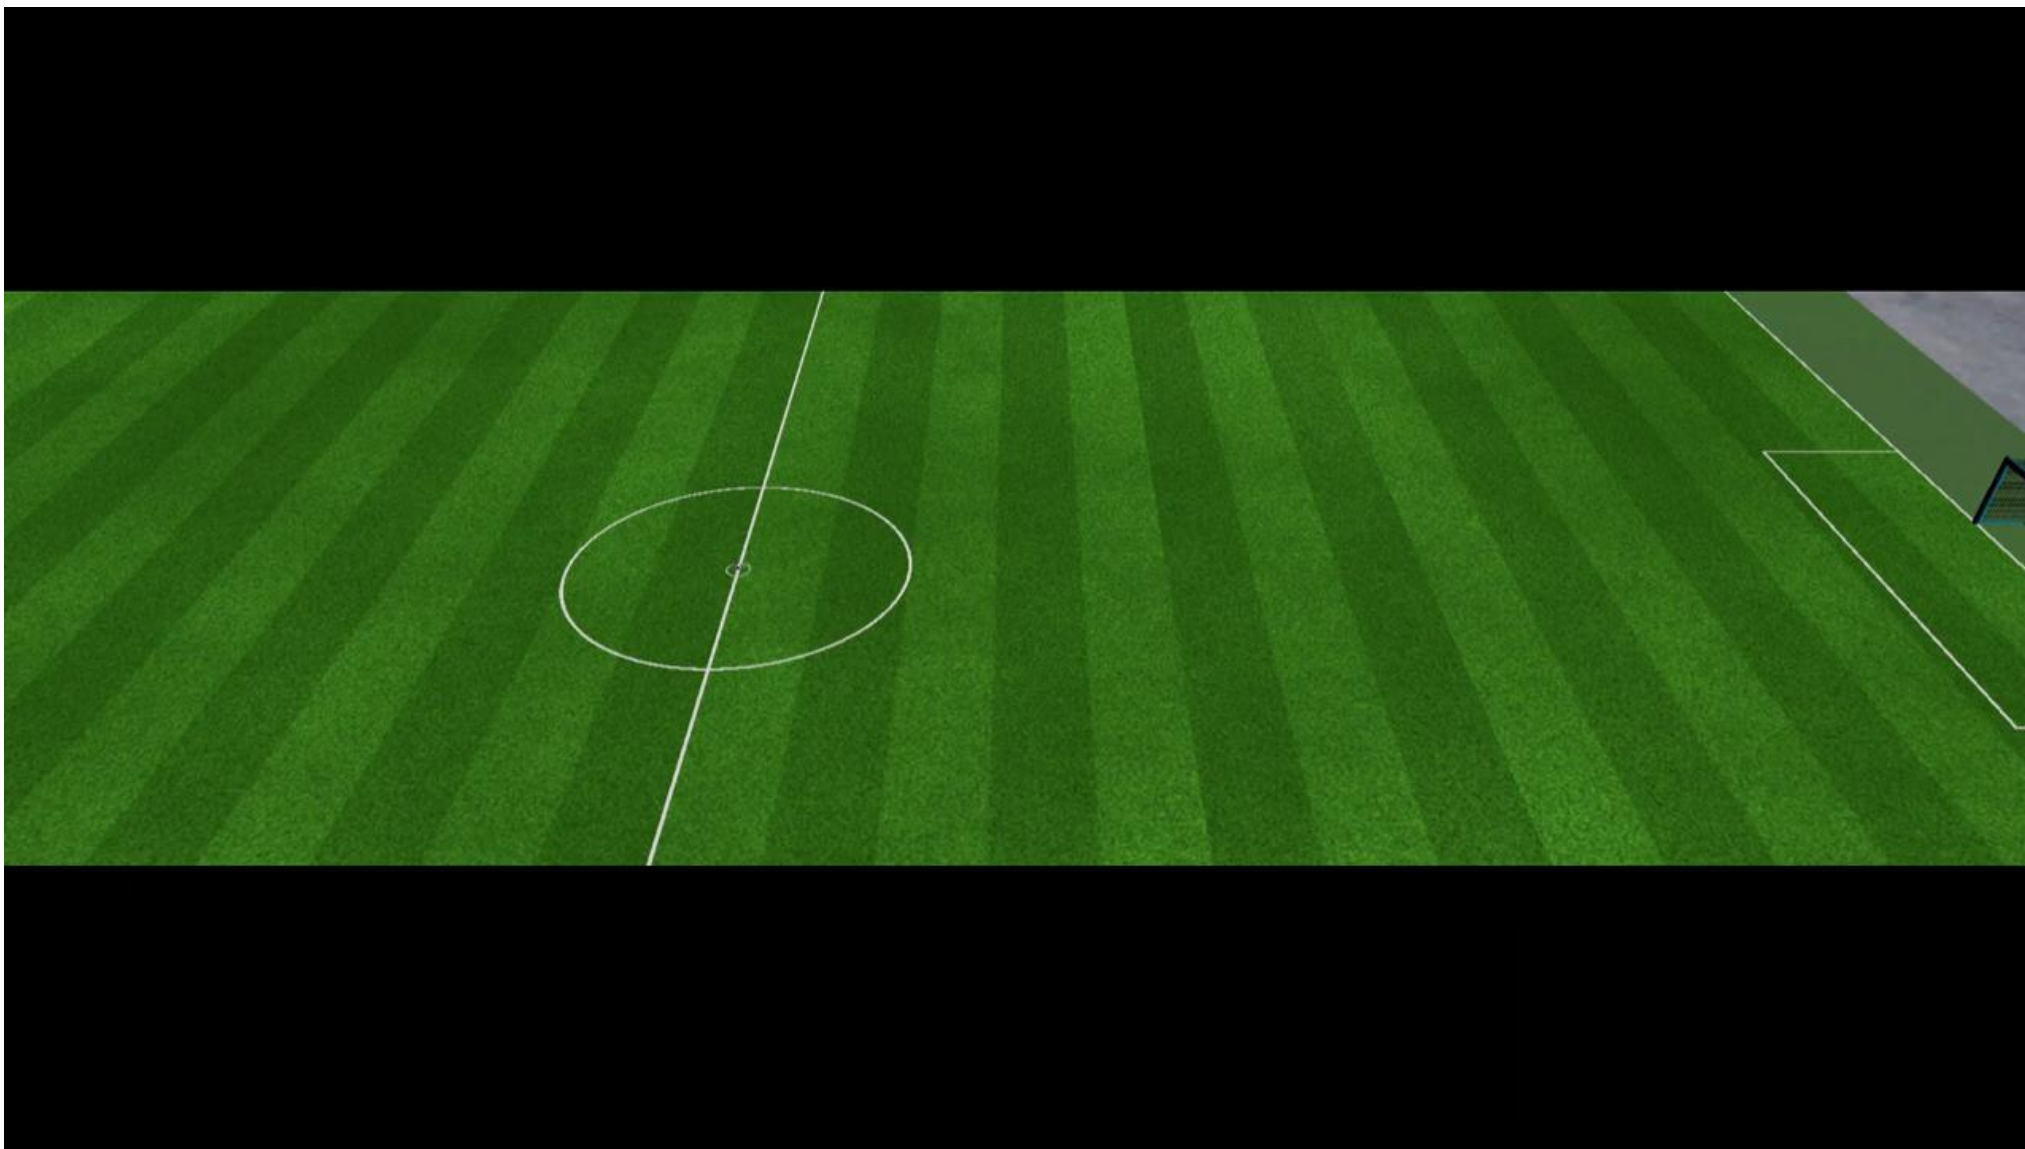
Optimization is performed by the Proximal Policy Optimization algorithm extended with Proximal Symmetry Loss

6D Pose Localization

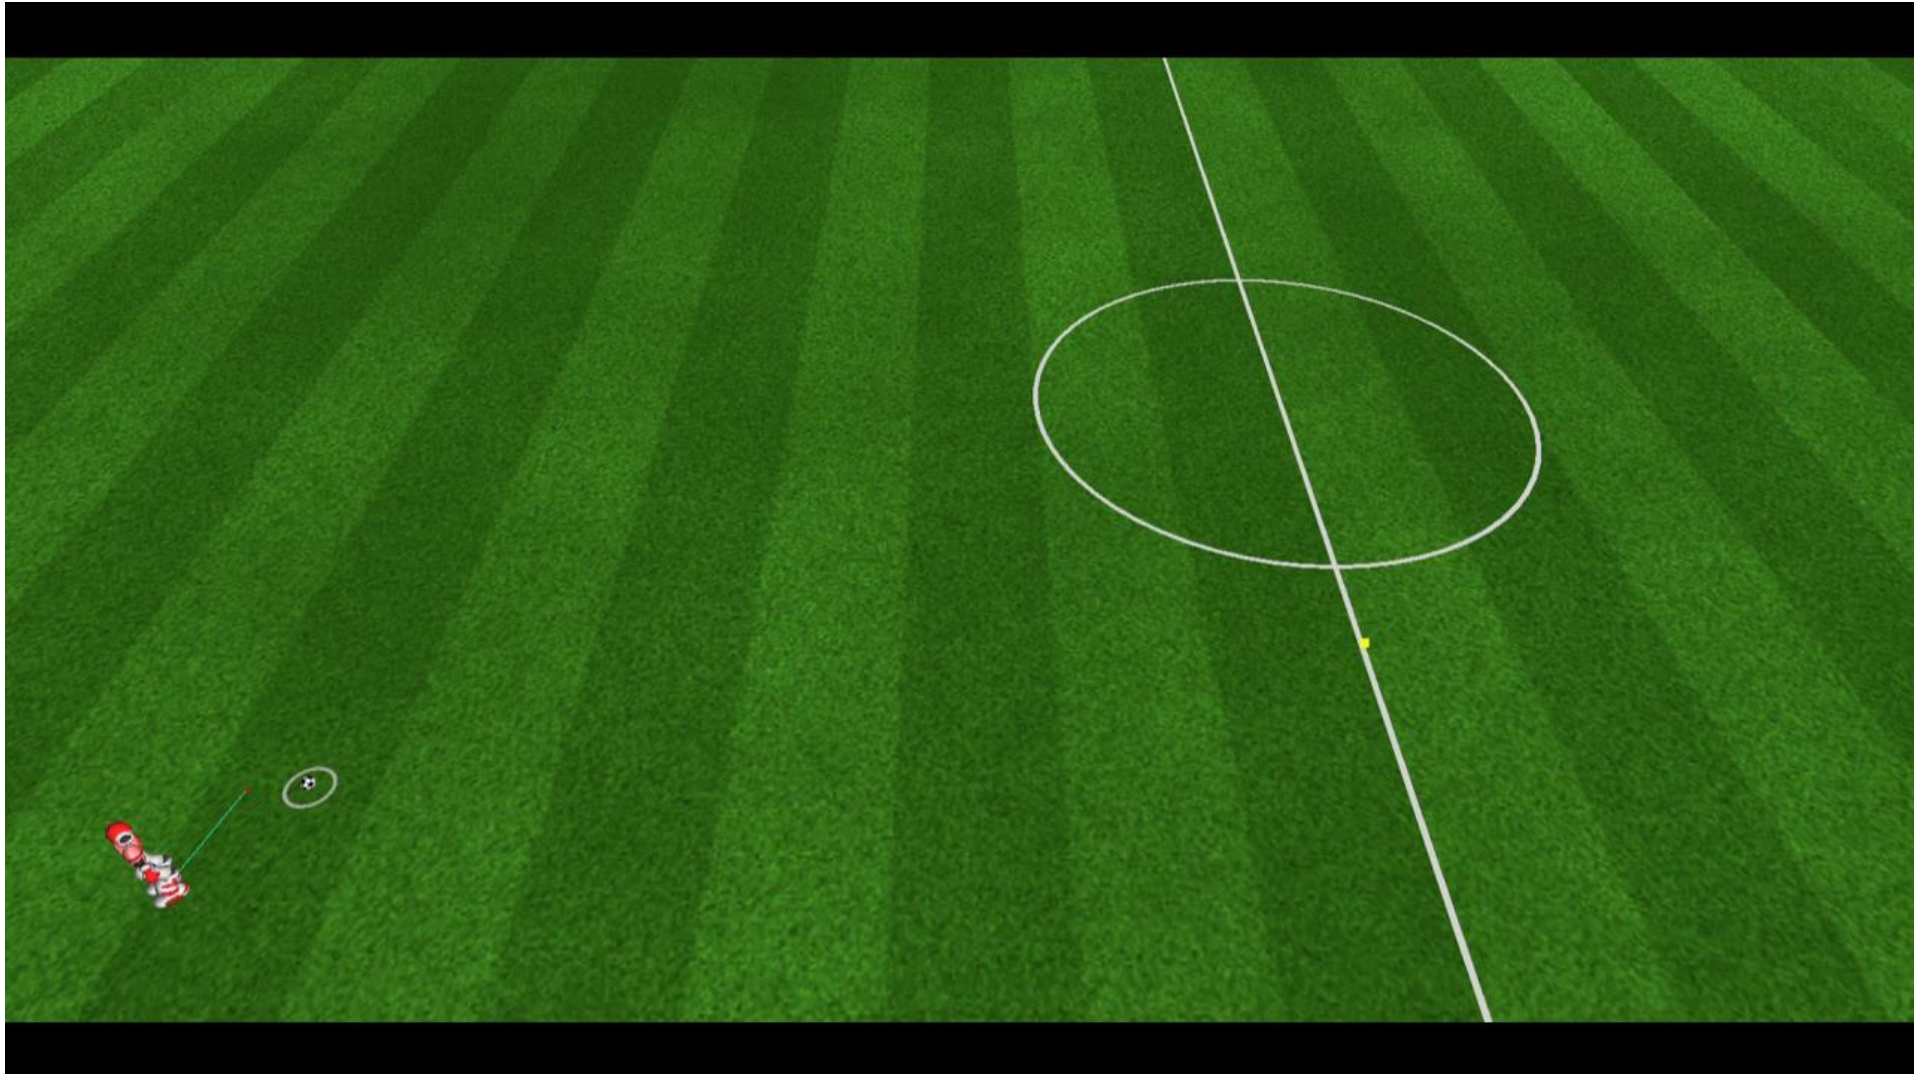


- **Robot localization algorithm**, estimates **6D pose** (position and orientation) of a humanoid
- Integrates **landmark detection**, **geometric transformations**, and **probabilistic optimization**
- Final pose refined using **Probabilistic Gradient Ascent**, which optimizes the position estimate based on likelihood distributions

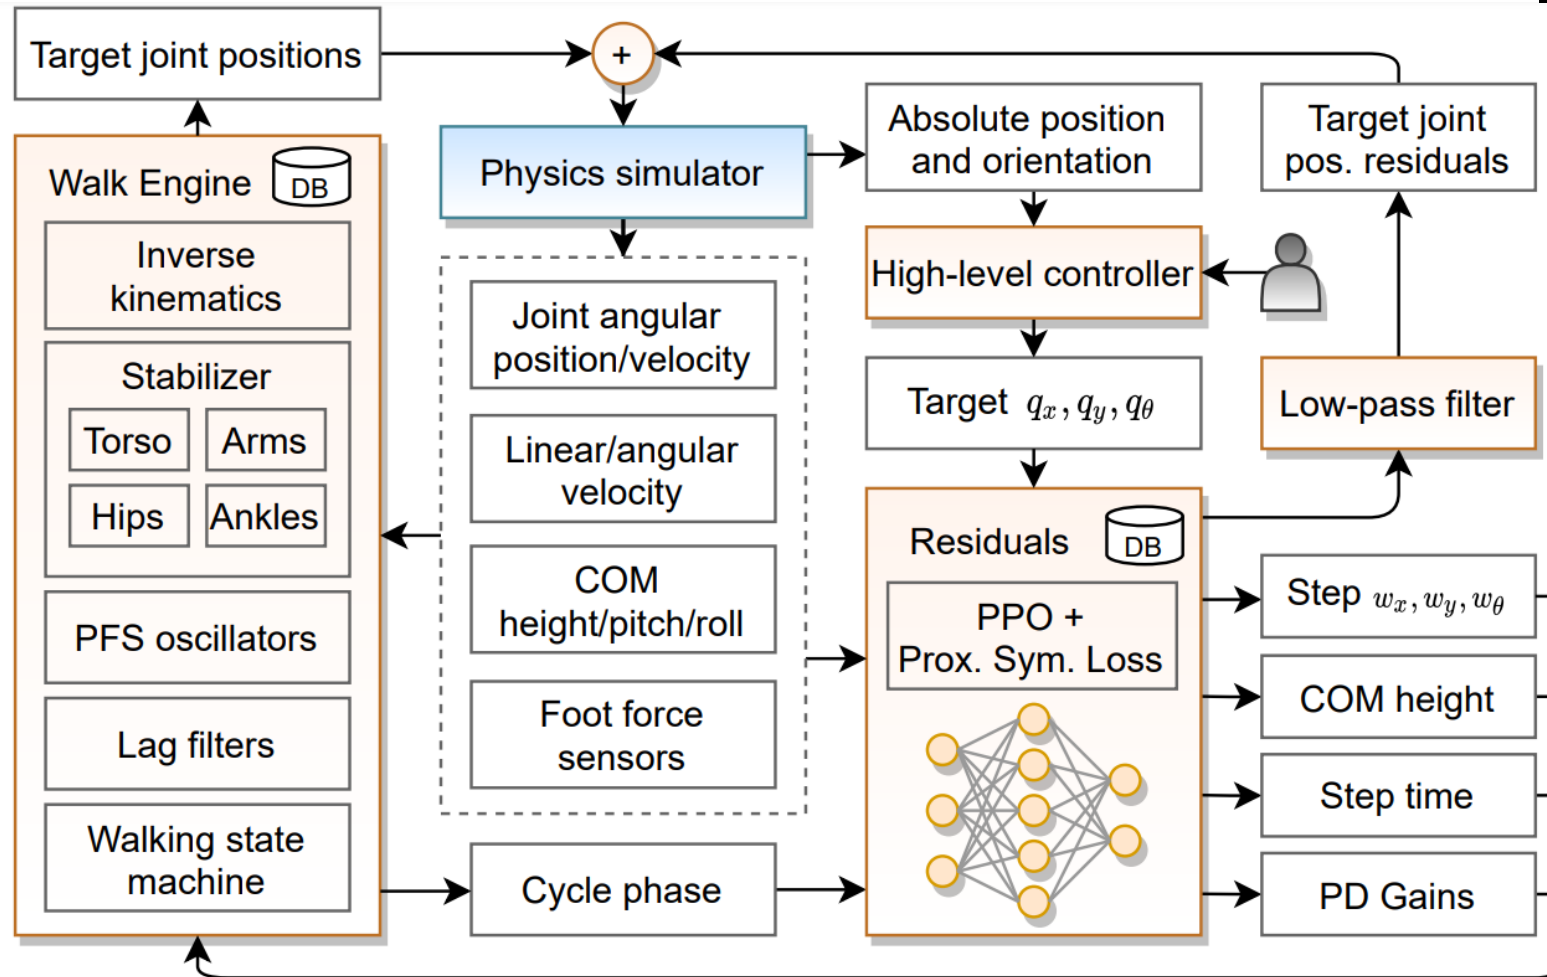
Learning to Kick Long



Learning to Control Kick



Learning to OmniWalk



DRL for Learning to Sprint



(Our Approach) FCPortugal



UT Austin Villa (3DSSL Champion)

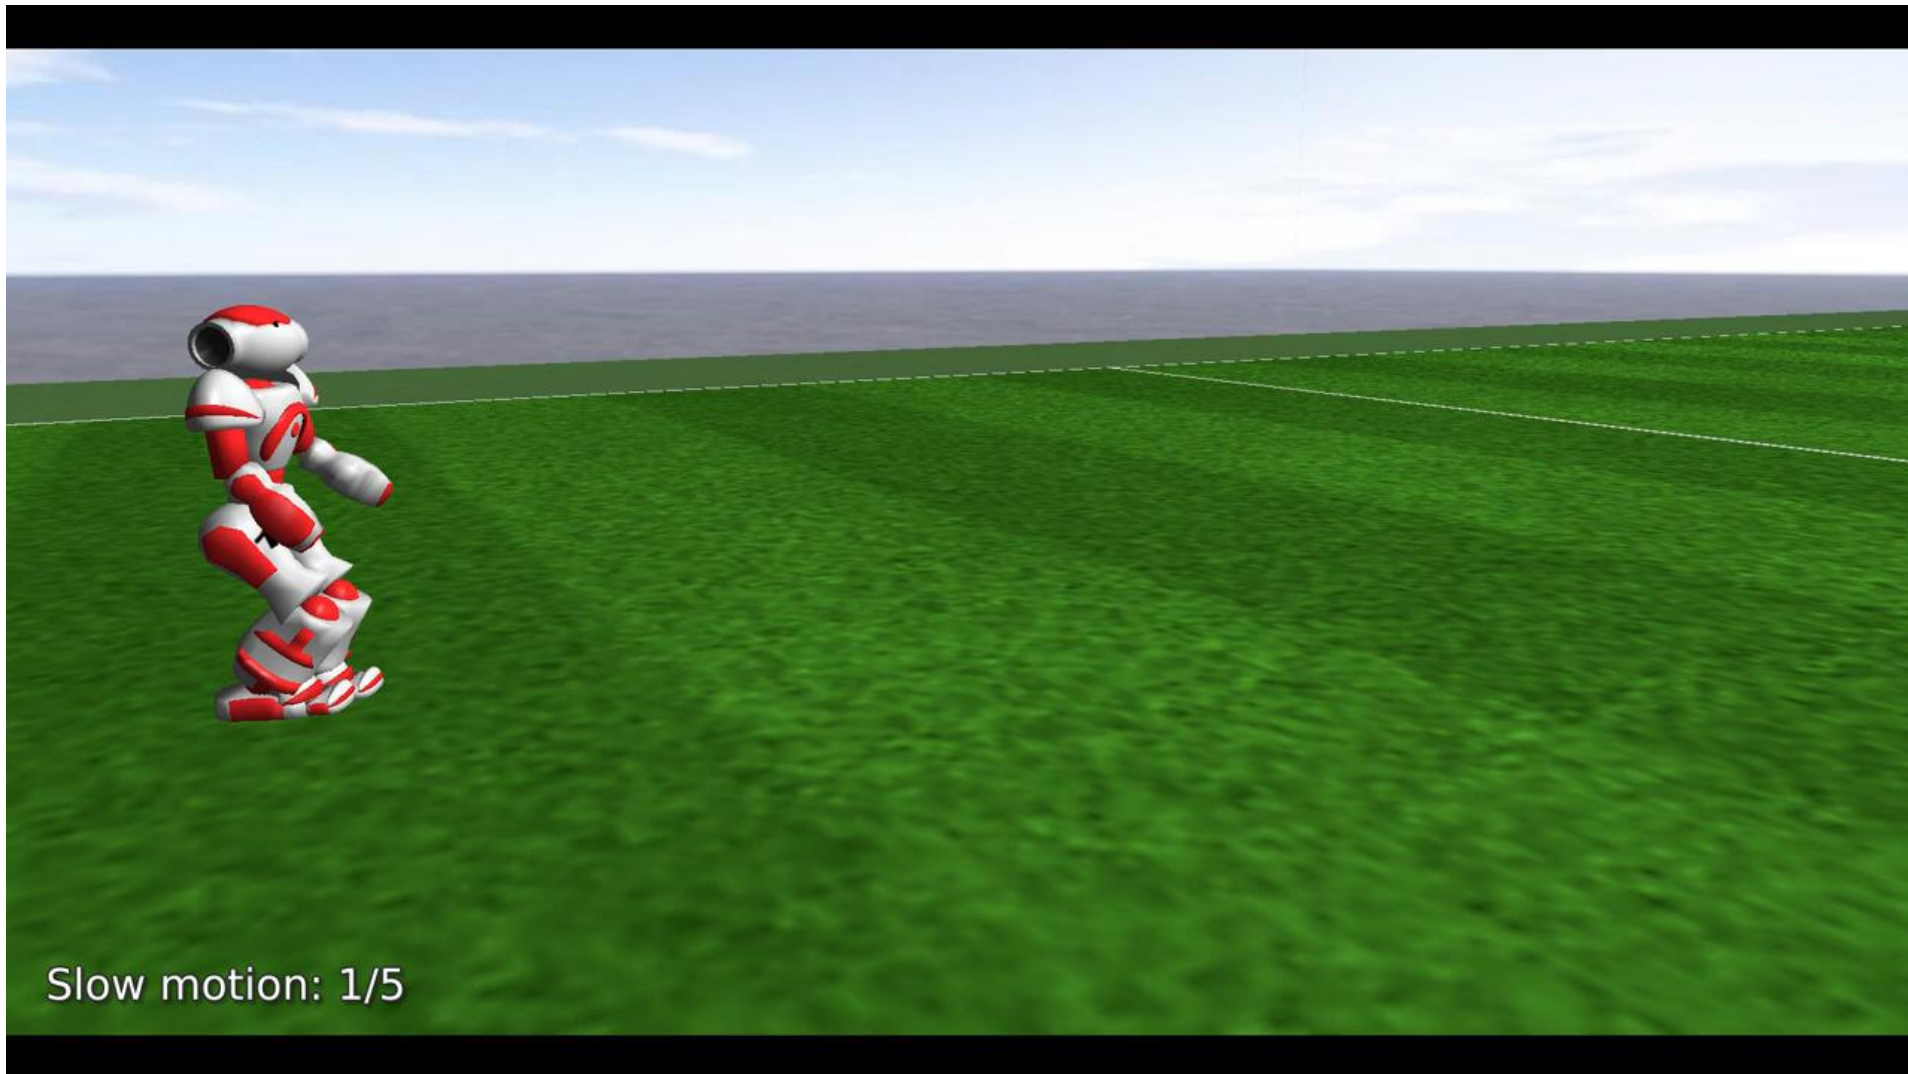


magmaOffenburg (3DSSL 2nd)

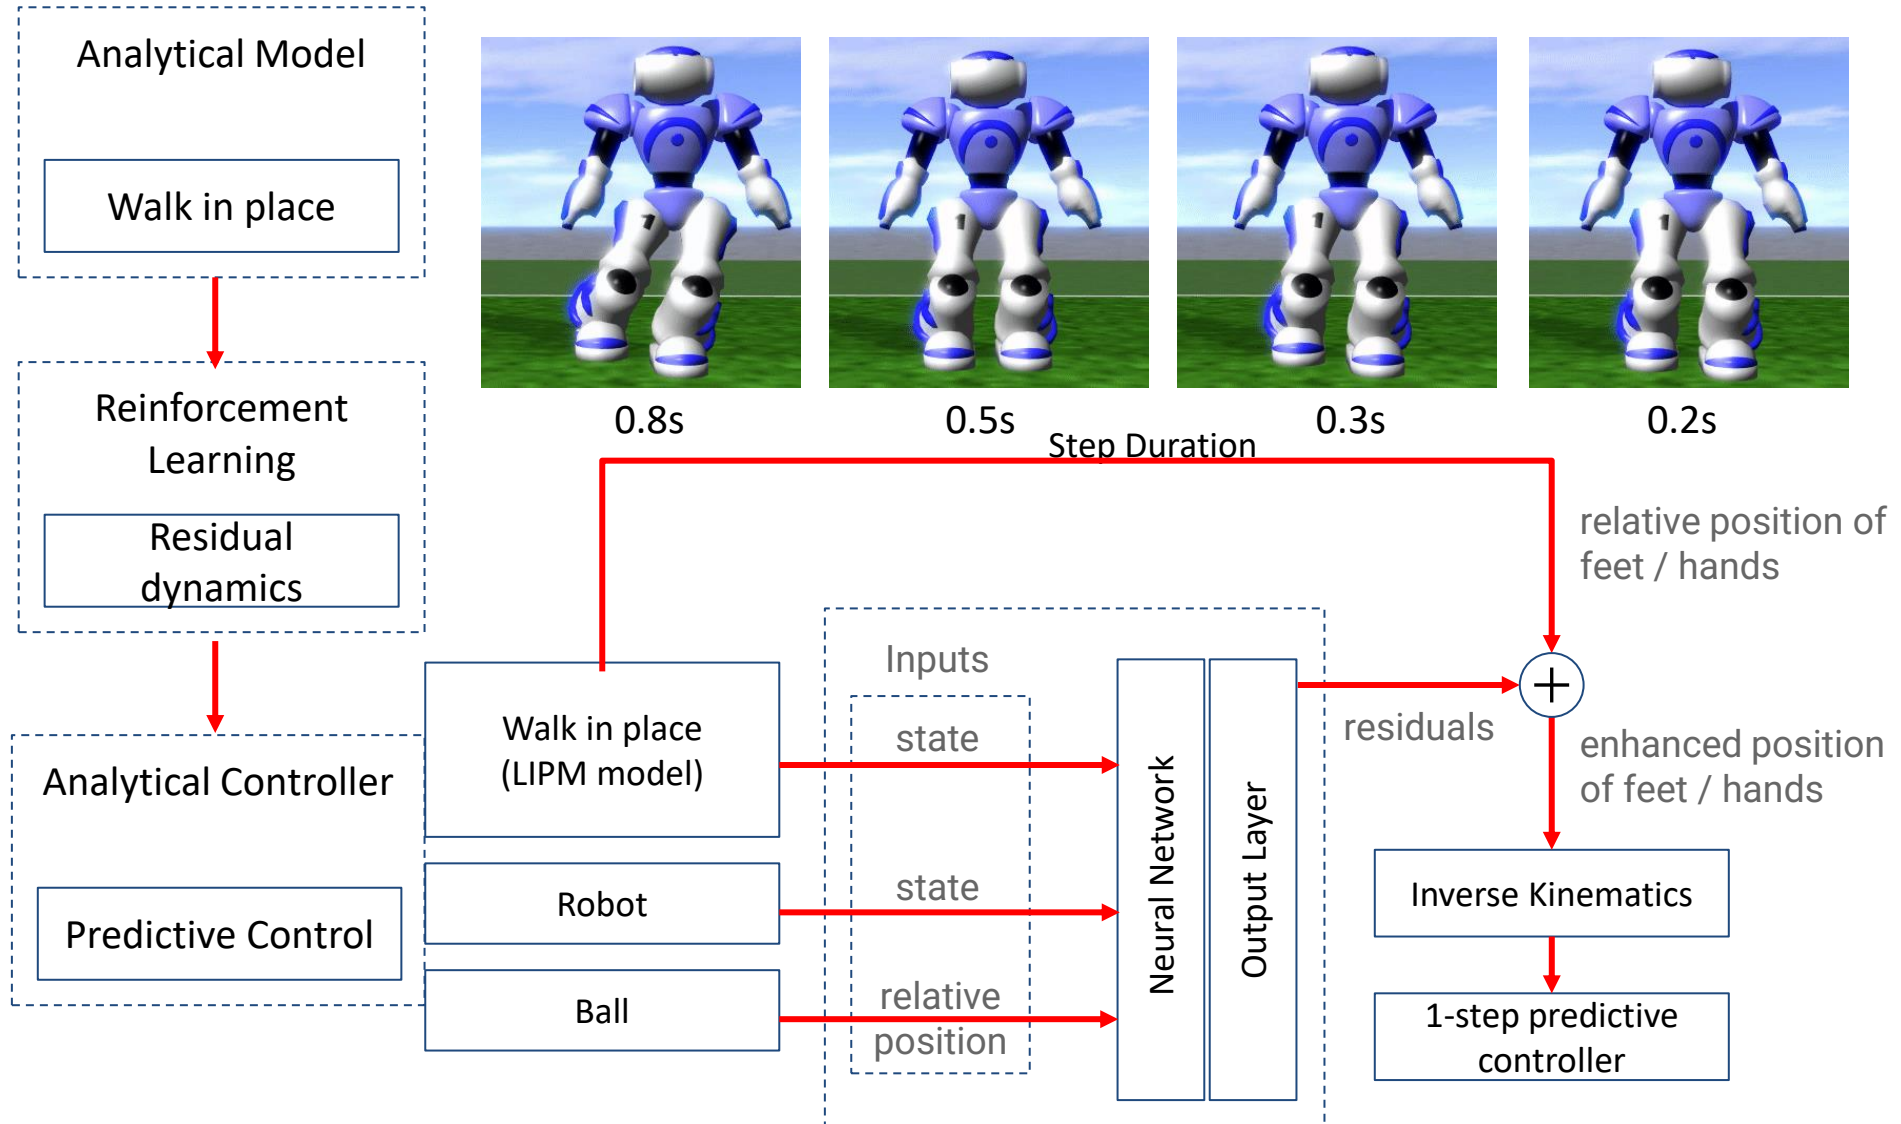
DRL for Learning to Sprint



DRL for Learning to Sprint and Kick



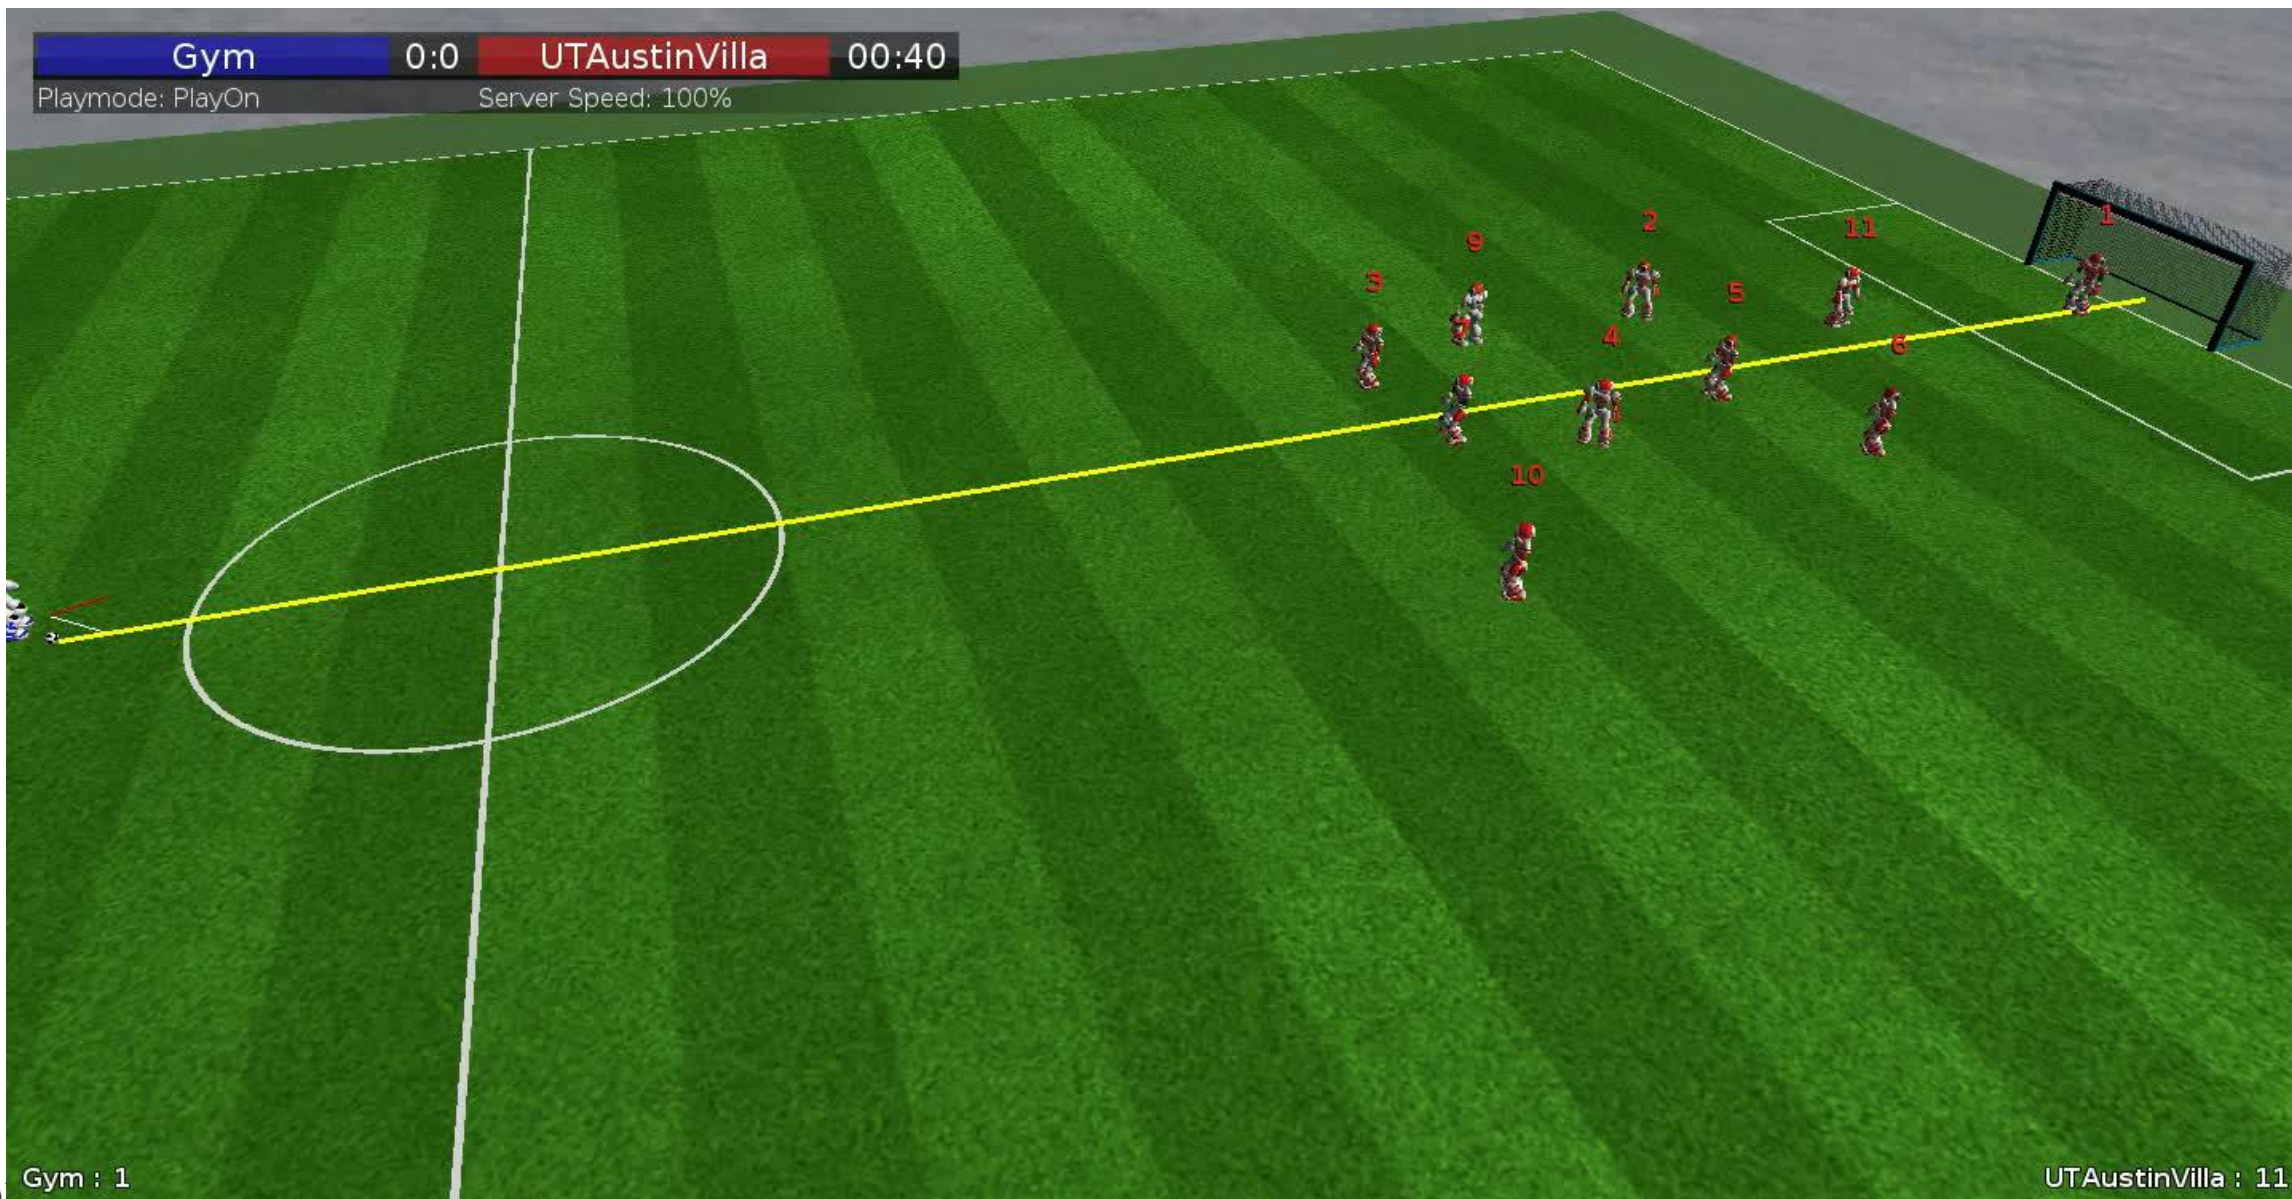
Learning to Dribble



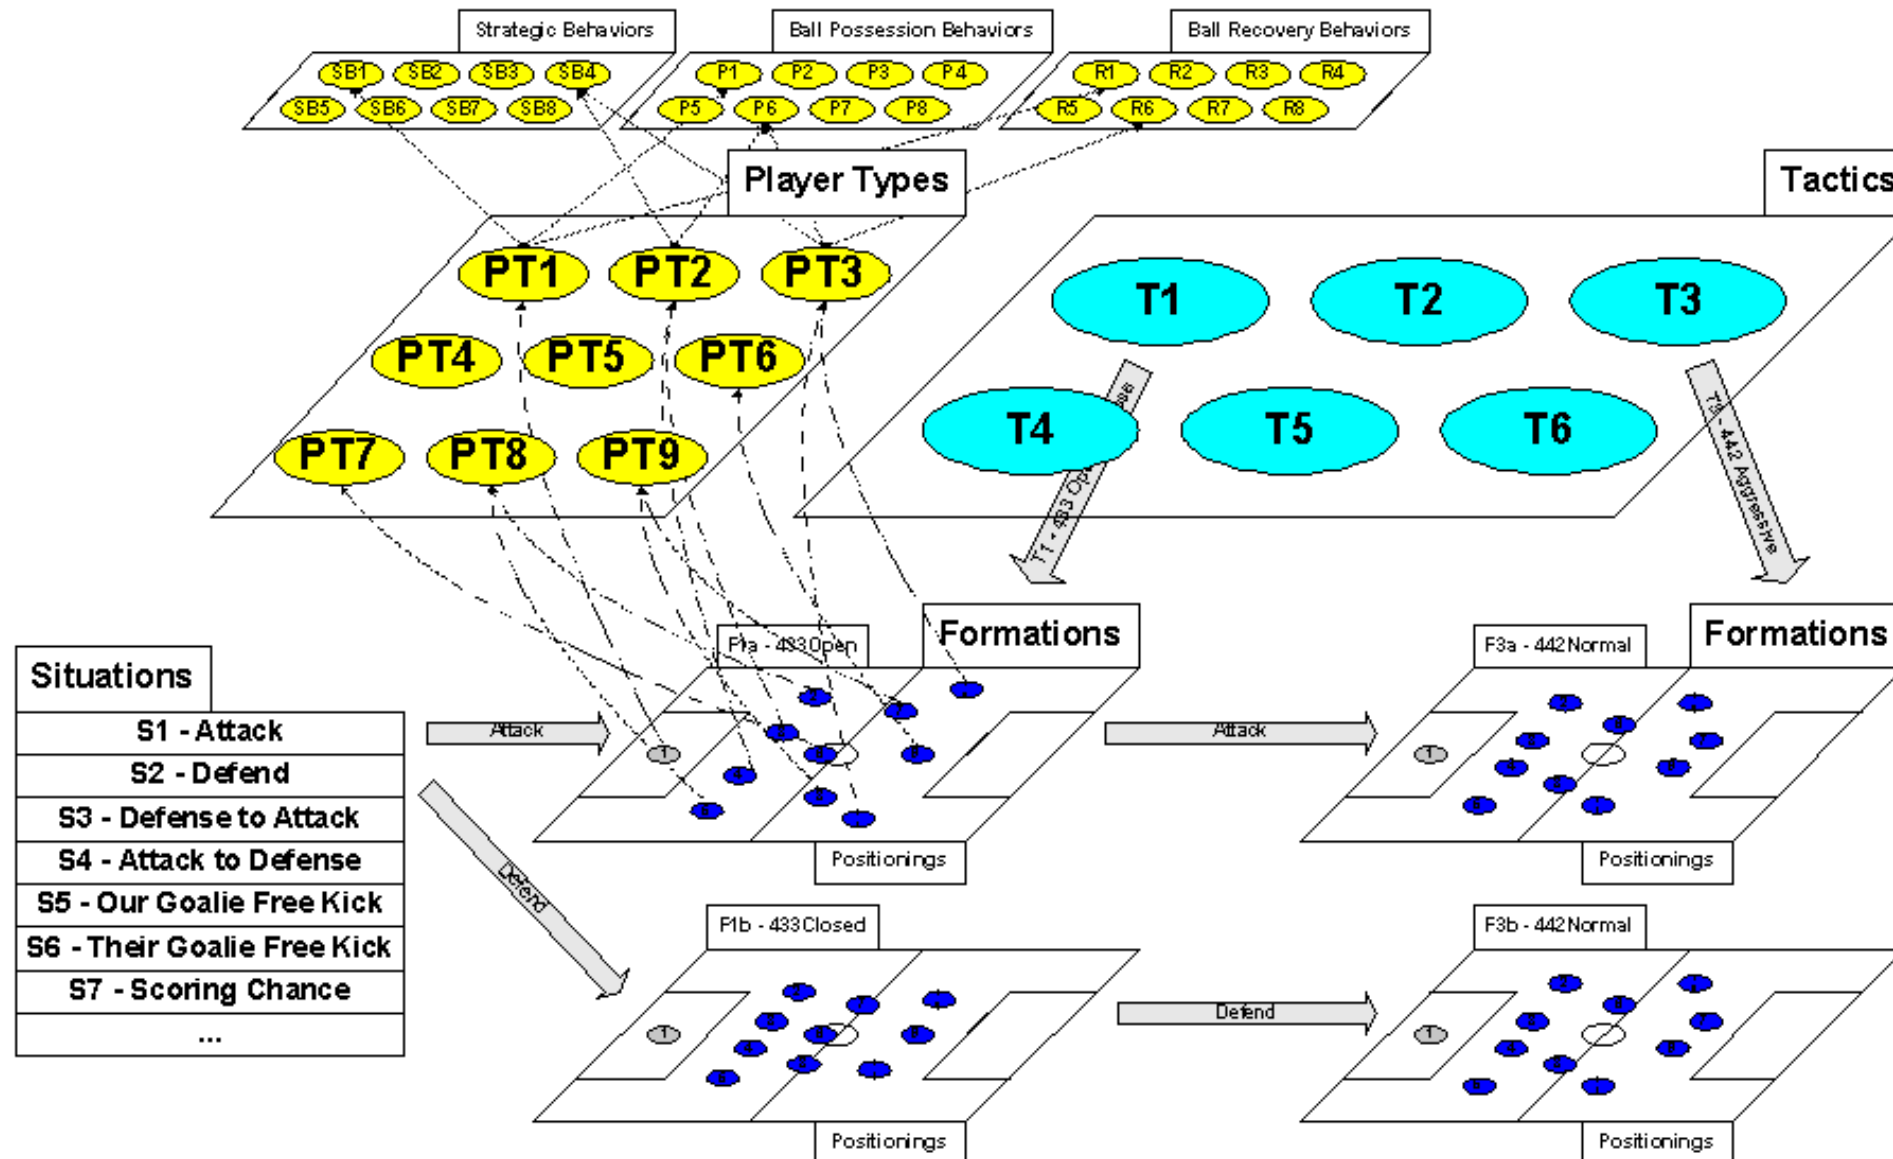
DRL for Learning to Dribble



DRL for Learning to Dribble



FC Portugal Team Strategy



RoboCup 2022 - Results



| Group A | magmaOffenburg | | BahiaRT | | Miracle3D | | FCPortugal | | HFUTEngine | |
|----------------|----------------|----|---------|---|-----------|---|------------|---|------------|---|
| magmaOffenburg | | | 10 | 0 | 7 | 0 | 0 | 0 | 3 | 0 |
| BahiaRT | 0 | 10 | | | 0 | 3 | 0 | 7 | 0 | 4 |
| Miracle3D | 0 | 7 | 3 | 0 | | | 0 | 6 | 0 | 3 |
| FCPortugal | 0 | 0 | 7 | 0 | 6 | 0 | | | 8 | 0 |
| HFUTEngine | 0 | 3 | 4 | 0 | 3 | 0 | 0 | 8 | | |

| Group C | FCPortugal | | Apollo3D | | HFUTEngine | | KgpKubs | | BahiaRT | |
|------------|------------|---|----------|---|------------|---|---------|---|---------|---|
| FCPortugal | | | 7 | 0 | 4 | 0 | 9 | 0 | 8 | 0 |
| Apollo3D | 0 | 7 | | | 1 | 2 | 9 | 0 | 6 | 0 |
| HFUTEngine | 0 | 4 | 2 | 1 | | | 5 | 0 | 3 | 0 |
| KgpKubs | 0 | 9 | 0 | 9 | 0 | 5 | | | 2 | 1 |
| BahiaRT | 0 | 8 | 0 | 6 | 0 | 3 | 1 | 2 | | |

| Group E | FCPortugal | | UTAV | | Apollo3D | | ITAndroids | |
|------------|------------|---|------|---|----------|---|------------|---|
| FCPortugal | | | 1 | 0 | 3 | 0 | 4 | 0 |
| UTAV | 0 | 1 | | | 3 | 2 | 8 | 0 |
| Apollo3D | 0 | 3 | 2 | 3 | | | 2 | 0 |
| ITAndroids | 0 | 4 | 0 | 8 | 0 | 2 | | |

| Group G | FCPortugal | | Miracle3D | | UTAV | | HFUTEngine | | magmaOffenburg | | Apollo3D | |
|----------------|------------|---|-----------|---|------|----|------------|---|----------------|---|----------|---|
| FCPortugal | | | 5 | 0 | 1 | 0 | 3 | 0 | 3 | 1 | 4 | 0 |
| Miracle3D | 0 | 5 | | | 0 | 10 | 0 | 4 | 0 | 9 | 0 | 3 |
| UTAV | 0 | 1 | 10 | 0 | | | 7 | 0 | 3 | 2 | 7 | 0 |
| HFUTEngine | 0 | 3 | 4 | 0 | 0 | 7 | | | 0 | 6 | 3 | 4 |
| magmaOffenburg | 1 | 3 | 9 | 0 | 2 | 3 | 6 | 0 | | | 7 | 1 |
| Apollo3D | 0 | 4 | 3 | 0 | 0 | 7 | 4 | 3 | 1 | 7 | | |

Semi Final: **FCPortugal 5 – Appollo3D 0**

Final: **FCPortugal 6 – magmaOffenburg 1**

RoboCup Best Goals



Rule Changes!

Sim3D Technical Committee

- **Rules Discussion**

- Limit Dribbling to:

- 5 seconds, and
 - 5 meters

Solution: Dribble Retrain!

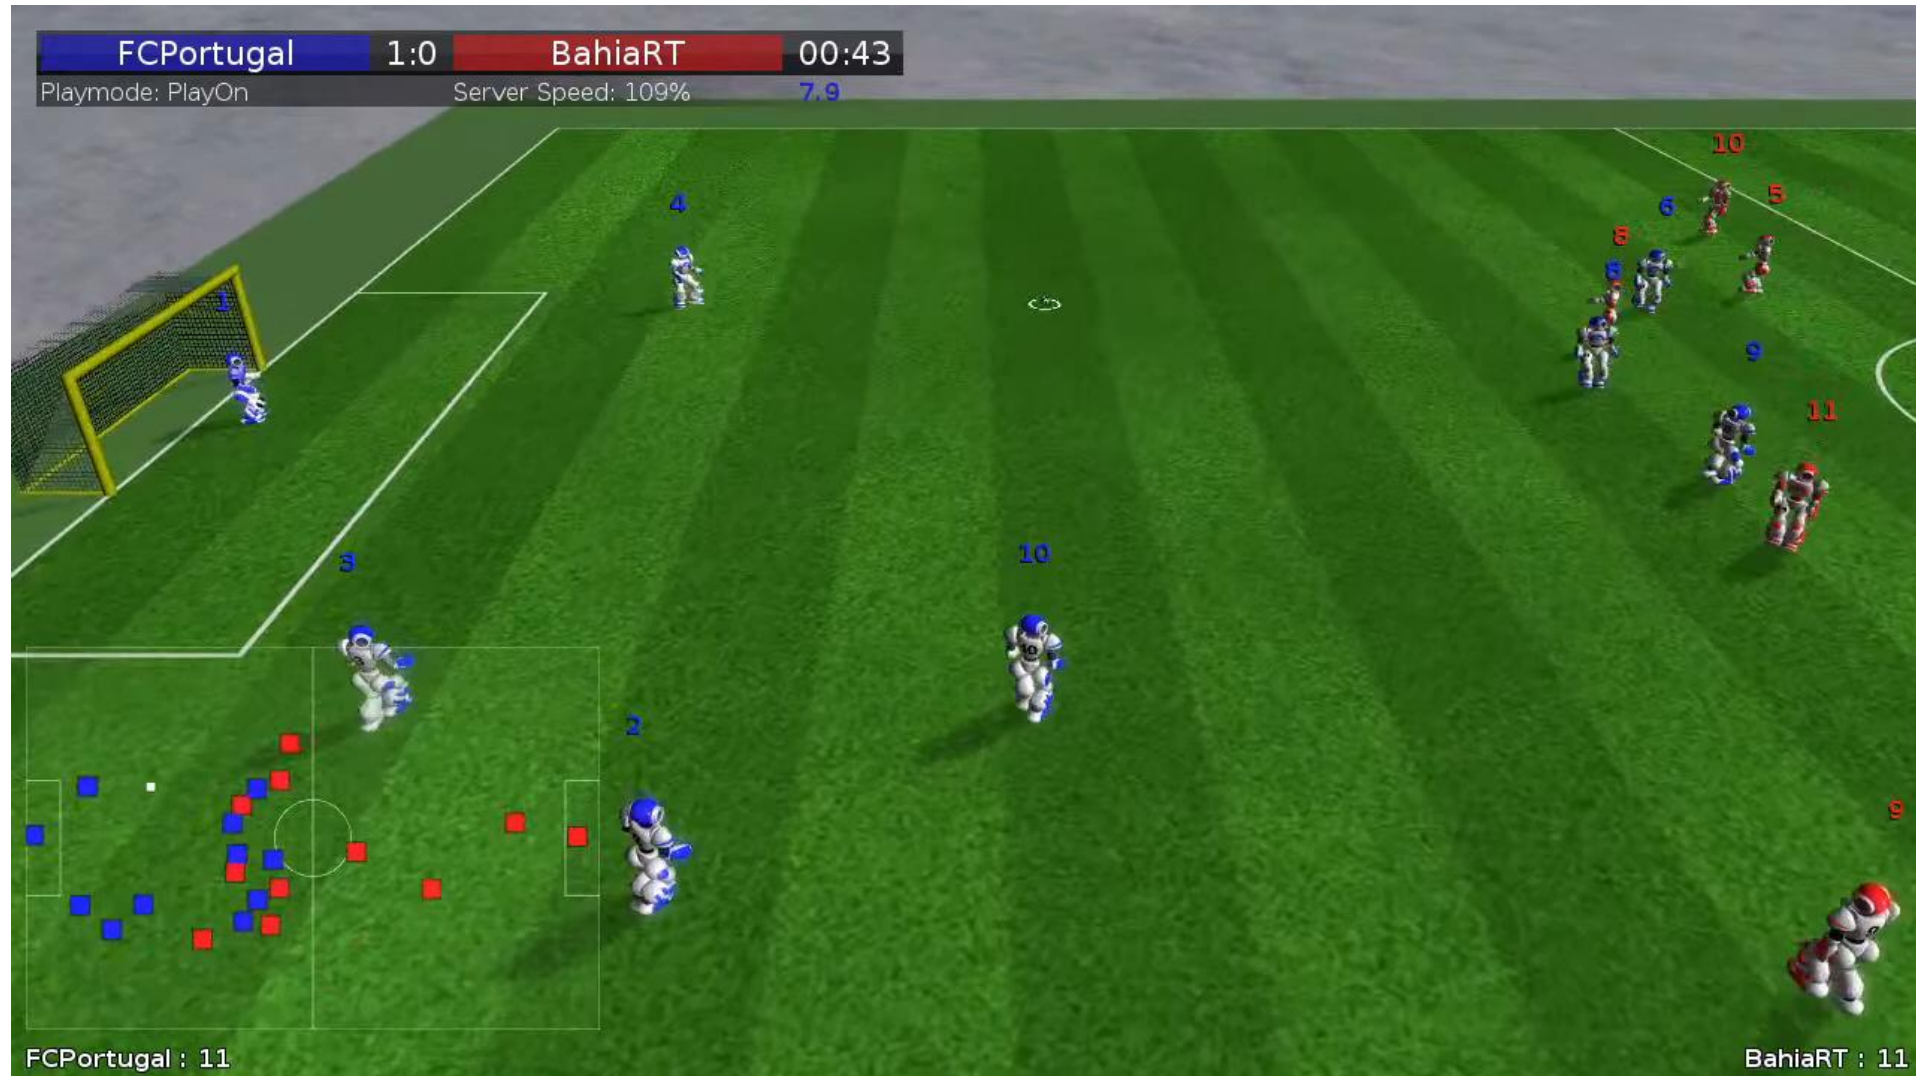
- **Final Rule**

- 2 weeks before the competition!
 - Ball Holding is forbidden!
 - No Dribbling?!

DRL for Learning again to Dribble



RoboCup Games



RoboCup Semi-Final



RoboCup Final



Code Release - <https://github.com/m-abr/FCPCodebase>

Skills

- Get Ups, Walk, Step (latest version)
- Dribbles
- Basic kick, and goalkeeper dive

Features

- Accurate localization: probabilistic 6D pose estimation algorithm and IMU
- Automatic communication: teammates to share location of players and ball
- Basics: common math ops, server communication, RoboViz drawings
- Custom A* pathfinding optimized for the soccer environment
- Easy integration of neural-network-based behaviors (via OpenAI Gym)
- Interactive demonstrations, tests and utilities showcasing key features of the team/agents
- Other modules: Relative/absolute position & orientation of every body part & joint through forward kinematics and vision, Inverse Kinematics, predictor for rolling ball position and velocity

Conclusions

- **AI and the New Discovery of the New World**
- **Strong Impact of AI in Research and Robotics**
- **New AI/GenAI, LBM's and DRL powered and trained Robots**
- **Approach to optimize robot performance in dynamic environments:
Accurate localization, efficient movement, and improved decision-making in RoboCup**
- **Python Code Base Release**
- **Huge Success in recent competitions – Tri-World Champions (2022, 2023, 2024)**

FC Portugal: Tri-World Champions in RoboCup 3D Humanoid Soccer Simulation

Luís Paulo Reis

lpreis@fe.up.pt

Associate Professor at FEUP - Faculty of Engineering of the University of Porto
Director of LIACC – Artificial Intelligence and Computer Science Laboratory, Univ. Porto
Member of Coordinating Commission of LASI – Intelligent Systems Associate Laboratory
President of the GA of APPIA – Portuguese Association for Artificial Intelligence

