*Slides completos, com anotações, animações e vídeos, encontram-se nesse link (Google Slides): https://docs.google.com/presentation/d/e/2PACX-1vTOssRXnPsr3QiT3cO_qejAHYxos1gl0XFEJ81-PPzHeD3umf4E257JMukDrmJzIDeL2taLx8zKVye6/pub?start=true&loop=false&delayms=30000

Bots autônomos em StarCraft II

Criando uma IA para o jogo usando o ambiente *PySC2*

Paulo Bruno de Sousa Serafim

Atlântico

Porto Alegre
2019

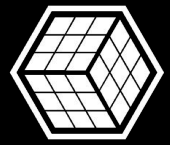

THE DEVELOPER'S CONFERENCE

paulo_serafim@atlantico.com.br
paulobruno@alu.ufc.br

Bots autônomos em StarCraft II

Criando uma IA para o jogo usando o ambiente *PySC2*

Nos últimos anos, o avanço nas técnicas de Inteligência Artificial reacenderam o interesse na disputa de humanos contra máquinas em jogos cada vez mais complexos. Em 2019, **dois dos melhores jogadores de StarCraft II foram derrotados por uma IA perdendo todas as partidas disputadas**. [...]

Bots autônomos em StarCraft II

Criando uma IA para o jogo usando o ambiente *PySC2*

Nos últimos anos, o avanço nas técnicas de Inteligência Artificial reacenderam o interesse na disputa de humanos contra máquinas em jogos cada vez mais complexos. Em 2019, **dois dos melhores jogadores de StarCraft II foram derrotados por uma IA perdendo todas as partidas disputadas.** [1]

BLOG POST
RESEARCH

30 OCT 2019

AlphaStar: Grandmaster level in StarCraft II using multi-agent reinforcement learning

Technology

DeepMind AI achieves Grandmaster status at Starcraft 2

By Leo Kelion
Technology desk editor

30 October 2019


StarCraft II

Lançado em 2010 pela Blizzard Entertainment

Real-time Strategy (RTS)

3 milhões de cópias vendidas em 1 mês [1]

Zerg

Terran

Protoss


PRODUCTION
Workers: 4/16
Workers: 0/3
5:57
Dusk Towers
BEST OF 3
0 <FCEO>PandaBear... 140 204 59/68 35 24
0 <Grav>EJK 351 300 66/61 38 26
Nydus Worm
Current Status

MICRO Gerenciamento

Controle individual de unidades

Ataques

Defesas

MACRO Gerenciamento

Economia

Gerenciamento de recursos

Construção de instalações

Atari

Go

StarCraft

Benchmarks of complexity	Increasing complexity →		
Information Type	Near-Perfect	Perfect	**Imperfect**
Players	Single player	**Multiplayer**	**Multiplayer**
Action Space	17	**361**	**~10^{26}**
Moves per game	100's	100's	**1000's**

DeepMind

Game	Board Size	State-Space Complexity	Year defeated
Tic Tac Toe	9	10^3	1952*
Connect 4	42	10^{13}	1995*
Backgammon	28	10^{20}	1979
Chess	64	10^{47}	1997
Go (19x19)	361	10^{170}	2015
Heads up NL Holdem	N/A	10^{180}	2017
StarCraft II	N/A	10^{1685}	???

@LIV_BOEREE

https://twitter.com/liv_boeree/status/1045025689710268421

BLOG POST
RESEARCH

24 JAN 2019

AlphaStar: Mastering the Real-Time Strategy Game StarCraft II

DeepMind

www.deepmind.com

	SUPPLY	MINERALS	GAS	WORKERS	ARMY	APM	PRODUCTION
0 AlphaStar	177 /200	945 +2015	758 +873	64	113	940	2
0 LiquidTLO	147 /172	335 +1595	442 +1030	61	86	1377	2 2

DeepMind | StarCraft II

DEMONSTRATION

TLO

Antes das partidas:
“Se eles já conseguirem me derrotar, isso seria incrível.” [2]

BLIZZARD ENTERTAINMENT

DeepMind | StarCraft II

DEMONSTRATION

TLO

ROUND	REPLAY
1.	ALPHASTAR WINS
2.	ALPHASTAR WINS
3.	ALPHASTAR WINS
4.	ALPHASTAR WINS
5.	ALPHASTAR WINS

SCORE TLO 0 - 5 ALPHASTAR

BLIZZARD ENTERTAINMENT

DeepMind | StarCraft II

DEMONSTRATION

TLO

“AlphaStar pega estratégias bem conhecidas e vira elas de cabeça pra baixo. O agente demonstrou estratégias que eu não tinha pensado antes, o que significa que pode haver novas maneiras de jogar StarCraft II que não explorei completamente ainda” [3]

BLIZZARD ENTERTAINMENT


PRODUCTION
DeepMind
StarCraft II
DEMONSTRATION
Workers: 0/16
Workers: 1/3
DeepMind
4:38
REPLAY
Catalyst LE
AlphaStar
37 /47
200 +671
76 +268
15
22
SUPPLY
MINERALS
GAS
WORKERS
ARMY
WEAPONS & ARMOR
LiquidMaNa
42 /55
145 +783
136 +291
26
15
ORACLE
0
69
31
Kills: 7

DeepMind | StarCraft II

DEMONSTRATION

GRZEGORZ 'MANA' KOMINCZ

Antes das partidas:

“Estou esperando um 5-0, sem perder nenhuma partida, mas eu acho que o objetivo realista seria 4-1 para mim.” [4]

BLIZZARD ENTERTAINMENT

DeepMind | StarCraft II

DEMONSTRATION

GRZEGORZ 'MANA' KOMINCZ

ROUND	REPLAY
1.	ALPHASTAR WINS
2.	ALPHASTAR WINS
3.	ALPHASTAR WINS
4.	ALPHASTAR WINS
5.	ALPHASTAR WINS

SCORE MANA 0 - 5 ALPHASTAR

BLIZZARD ENTERTAINMENT

DeepMind | StarCraft II

DEMONSTRATION

GRZEGORZ 'MANA' KOMINCZ

“Eu percebi o quanto meu gameplay depende de forçar erros e ser capaz de explorar reações humanas, então isso [as partidas contra AlphaStar] colocou o jogo sob uma luz totalmente nova para mim. Estamos todos muito animados para ver o que vem a seguir” [5]

BLIZZARD ENTERTAINMENT

JANUARY 24

DeepMind's AlphaStar AI wins 10-1 against professional StarCraft II players

Abner Li - Jan. 24th 2019 12:50 pm PT @technacity

DeepMind AI AlphaStar goes 10-1 against top 'StarCraft II' pros

The AI beat 'StarCraft' pros TLO and MaNa thanks to more than 200 years worth of game knowledge.

AJ Dellinger, @ajdell
01.24.19 in Robots

24
Comments

2743
Shares

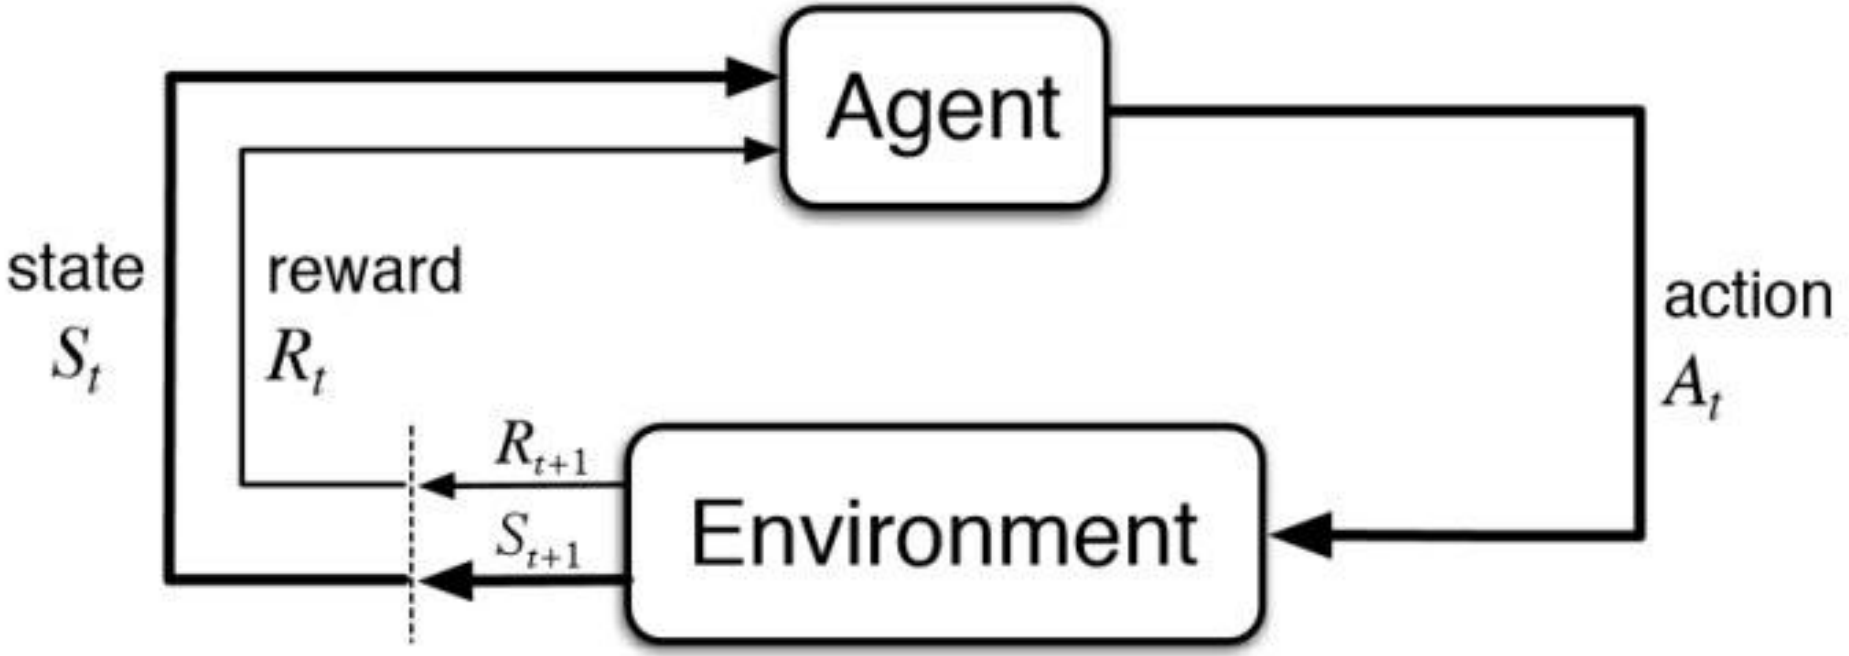

Sutton and Barto. Reinforcement Learning. 2018.

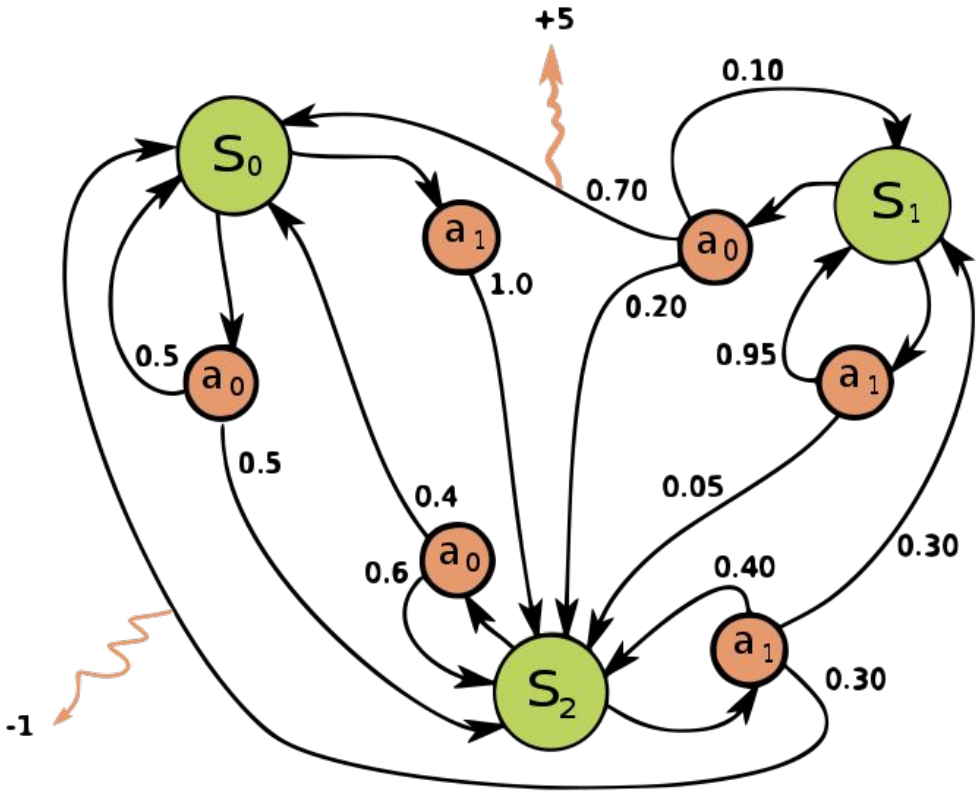

https://en.wikipedia.org/wiki/Markov_decision_process

PROCESSO DE MARKOV

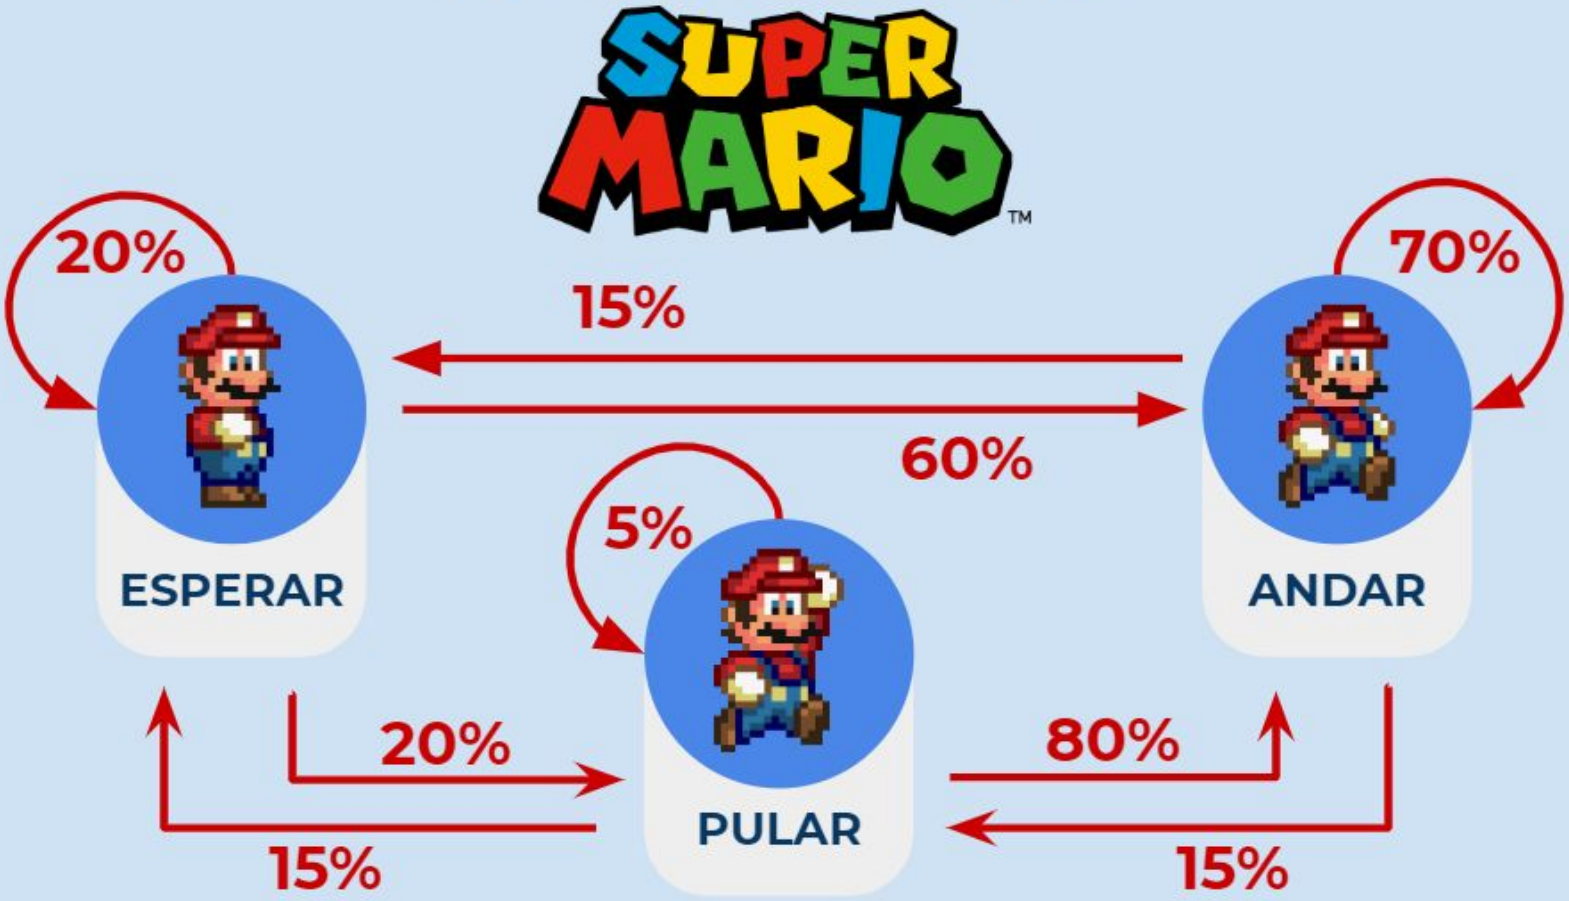

https://paulovasconcellos.com.br/explicando-deep-reinforcement-learning-com-super-mario-ao-inv%C3%A9s-de-matem%C3%A1tica-4c77392cc733

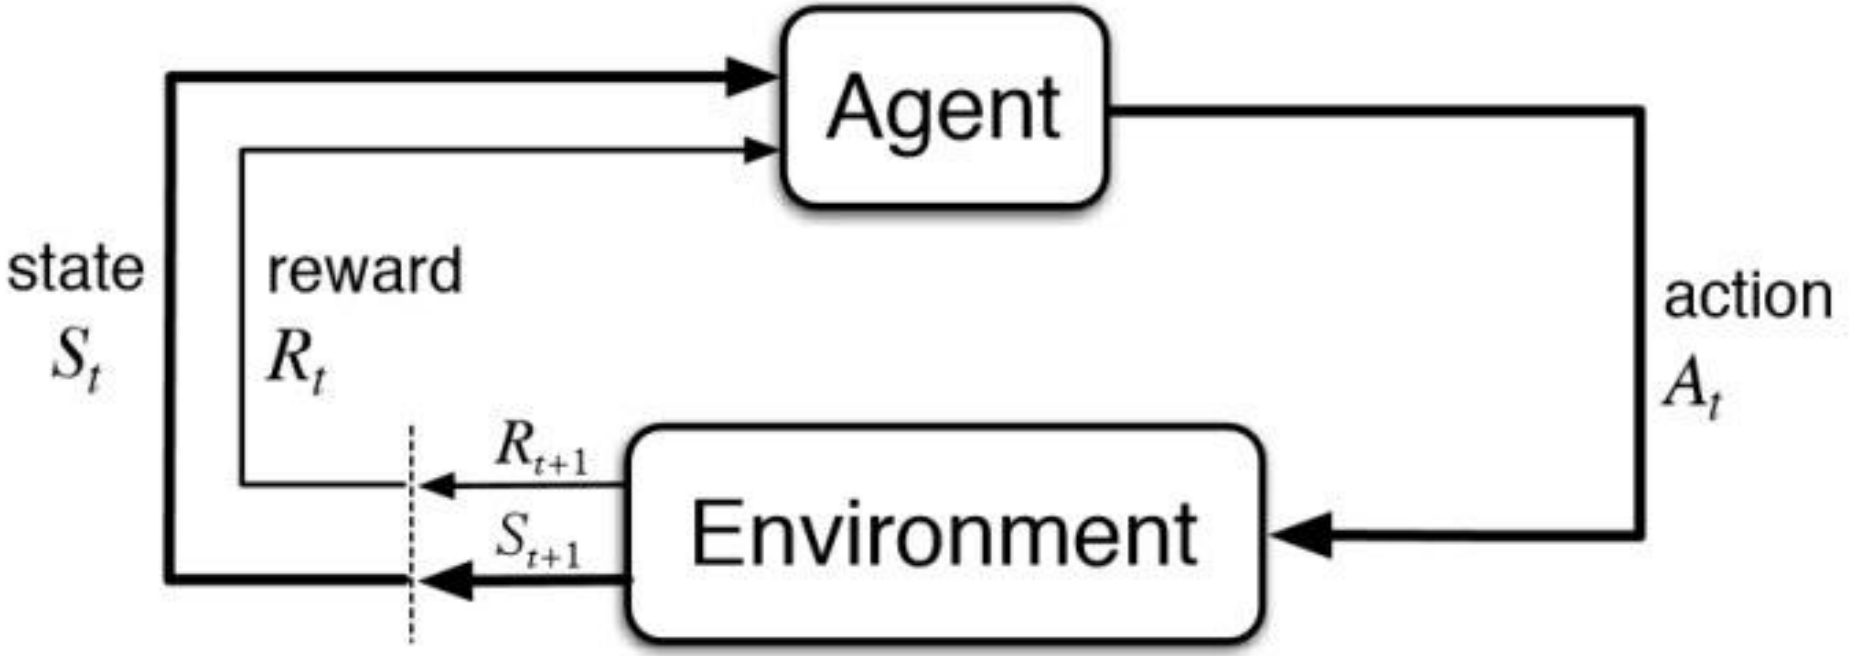

Sutton and Barto. Reinforcement Learning. 2018.

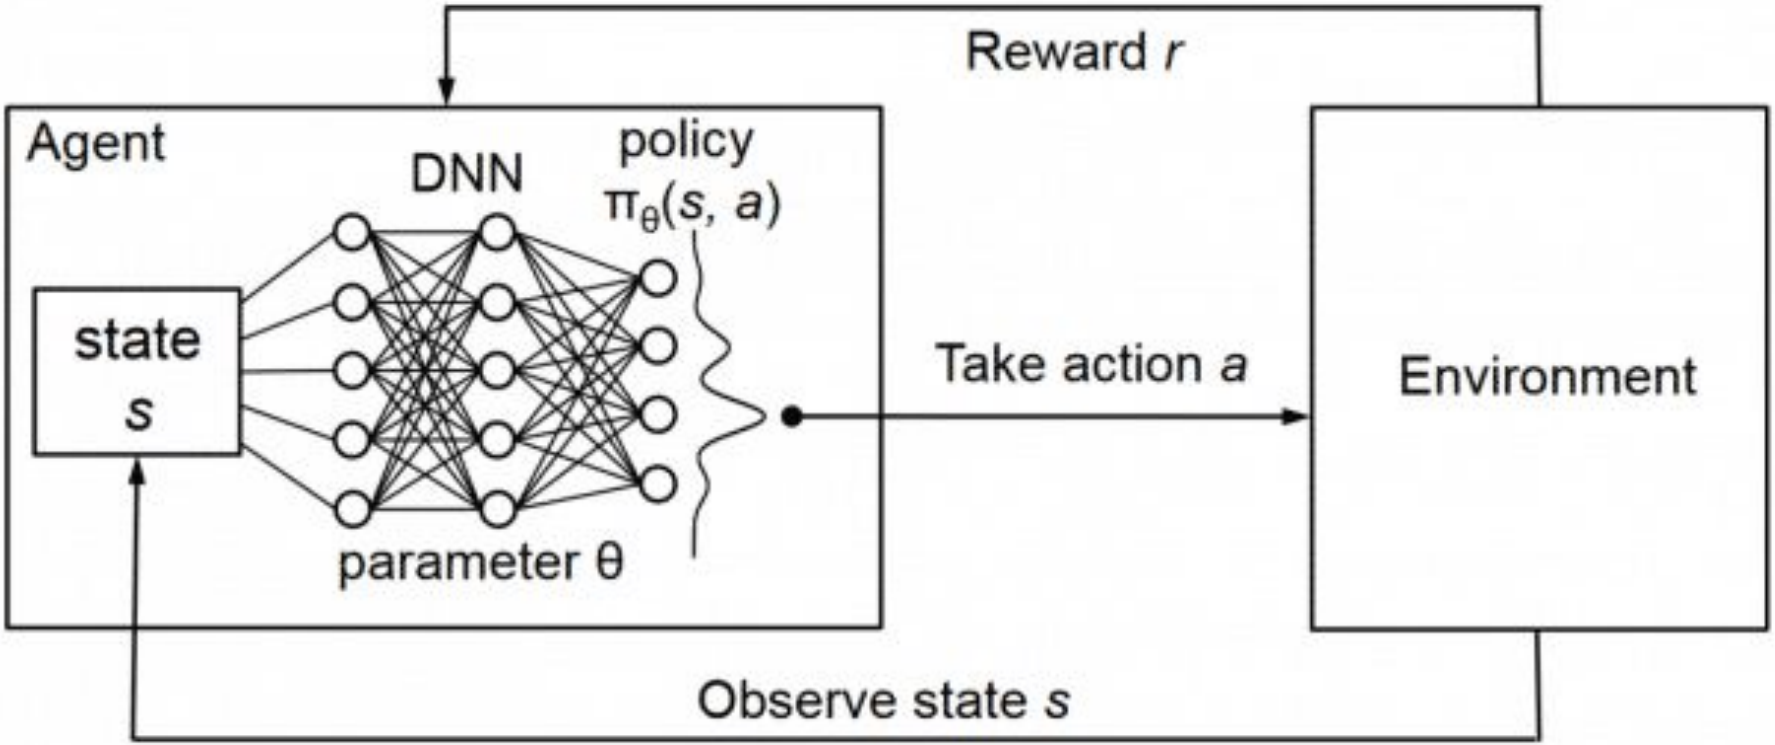

http://people.csail.mit.edu/hongzi/content/publications/DeepRM-HotNets16.pdf

Playing Atari with Deep Reinforcement Learning

Volodymyr Mnih **Koray Kavukcuoglu** **David Silver** **Alex Graves** **Ioannis Antonoglou**

Daan Wierstra **Martin Riedmiller**

DeepMind Technologies

{vlad,koray,david,alex.graves,ioannis,daan,martin.riedmiller} @ deepmind.com

Abstract

We present the first deep learning model to successfully learn control policies directly from high-dimensional sensory input using reinforcement learning. The model is a convolutional neural network, trained with a variant of Q-learning, whose input is raw pixels and whose output is a value function estimating future rewards. We apply our method to seven Atari 2600 games from the Arcade Learning Environment, with no adjustment of the architecture or learning algorithm. We find that it outperforms all previous approaches on six of the games and surpasses a human expert on three of them.

https://youtu.be/TmPfTpjtdgg

LETTER

doi:10.1038/nature14236

Human-level control through deep reinforcement learning

Volodymyr Mnih[1]*, Koray Kavukcuoglu[1]*, David Silver[1]*, Andrei A. Rusu[1], Joel Veness[1], Marc G. Bellemare[1], Alex Graves[1], Martin Riedmiller[1], Andreas K. Fidjeland[1], Georg Ostrovski[1], Stig Petersen[1], Charles Beattie[1], Amir Sadik[1], Ioannis Antonoglou[1], Helen King[1], Dharshan Kumaran[1], Daan Wierstra[1], Shane Legg[1] & Demis Hassabis[1]

The theory of reinforcement learning provides a normative account[1], deeply rooted in psychological[2] and neuroscientific[3] perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. Remarkably, humans and other animals seem to solve this problem through a harmonious combination of reinforcement learning and hierarchical sensory processing systems[4,5], the former evidenced by a wealth of neural data revealing notable parallels between the phasic signals emitted by dopaminergic neurons and temporal difference reinforcement learning algorithms[3]. While reinforcement learning agents have achieved some successes in a variety of domains[6–8], their applicability has previously been limited to domains in which useful features can be handcrafted, or to domains with fully observed, low-dimensional state spaces. Here we use recent advances in training deep neural networks[9–11] to

agent is to select actions in a fashion that maximizes cumulative future reward. More formally, we use a deep convolutional neural network to approximate the optimal action-value function

$$Q^*(s,a) = \max_{\pi} \mathbb{E}\left[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \ldots \mid s_t = s,\, a_t = a,\, \pi\right],$$

which is the maximum sum of rewards r_t discounted by γ at each time-step t, achievable by a behaviour policy $\pi = P(a|s)$, after making an observation (s) and taking an action (a) (see Methods)[19].

Reinforcement learning is known to be unstable or even to diverge when a nonlinear function approximator such as a neural network is used to represent the action-value (also known as Q) function[20]. This instability has several causes: the correlations present in the sequence of observations, the fact that small updates to Q may significantly change the policy and therefore change the data distribution, and the correlations between the action-values (Q) and the target values $r + \gamma \max_{a'} Q(s', a')$. We address these instabilities with a novel variant of Q-learning, which uses two key ideas. First, we used a biologically inspired mechanism termed experience replay[21–23] that randomizes over the data, thereby

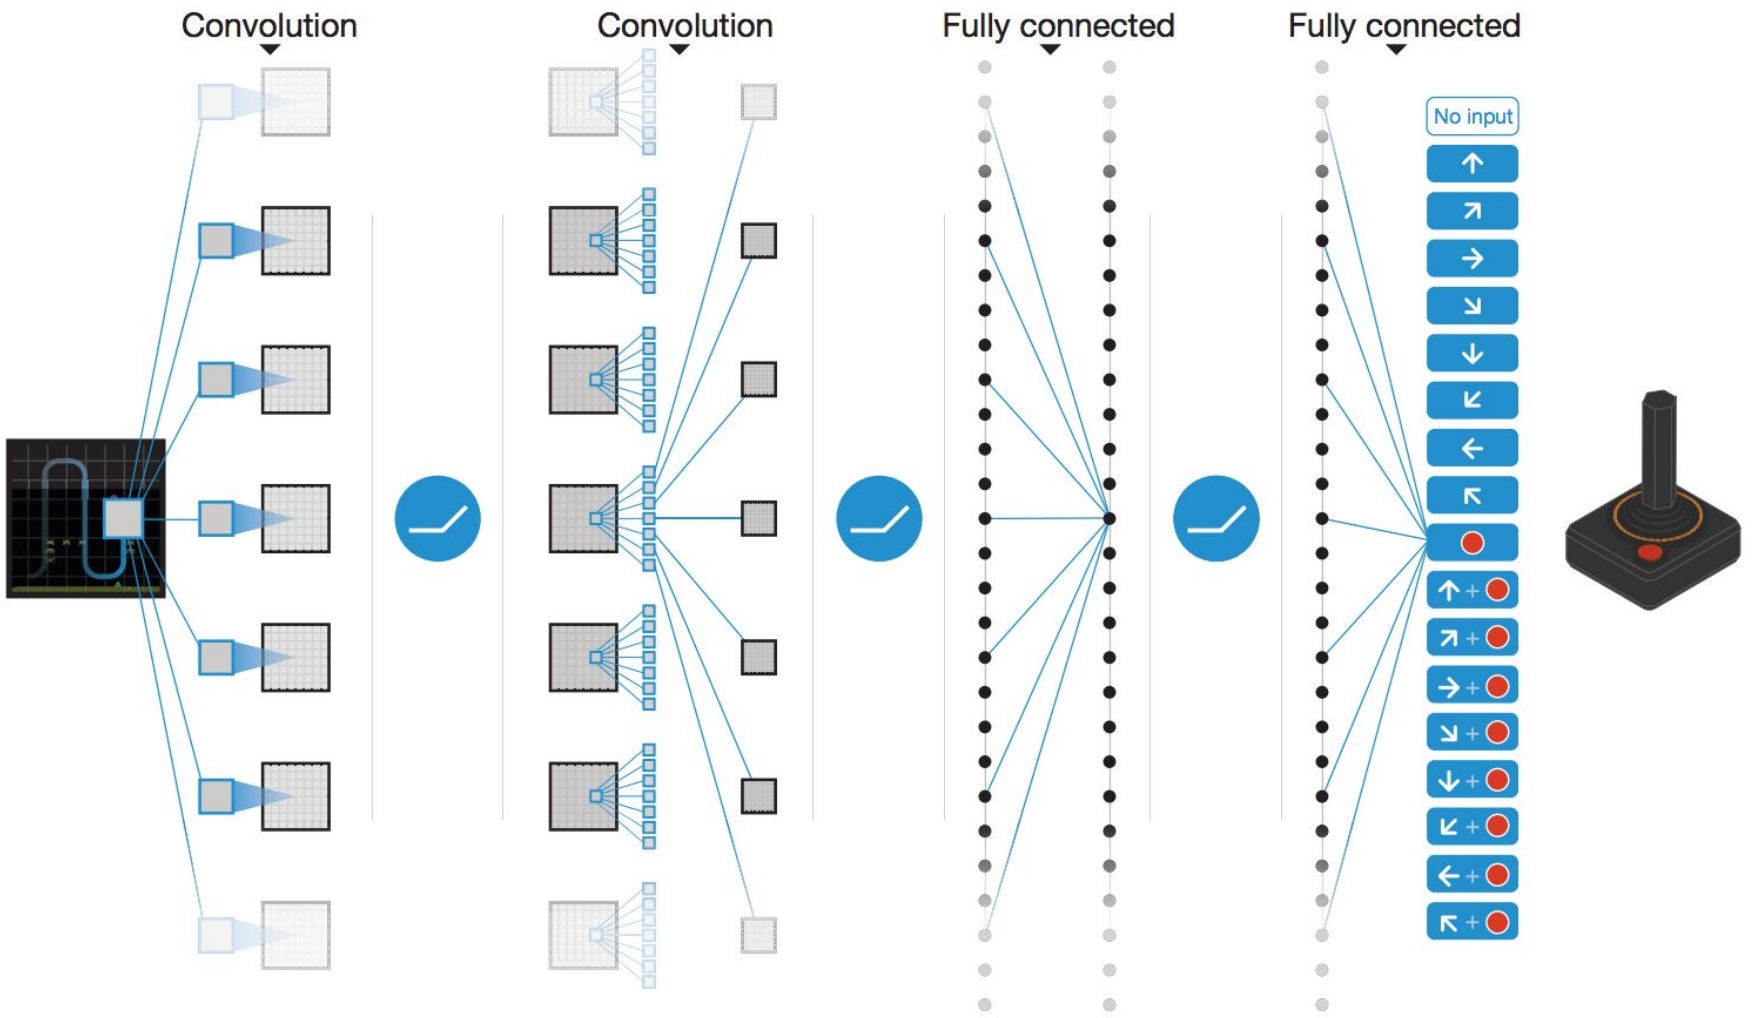
Convolution
Convolution
Fully connected
Fully connected
No input

AlphaGo *4* vs *1* Lee Sedol

Raw Observations

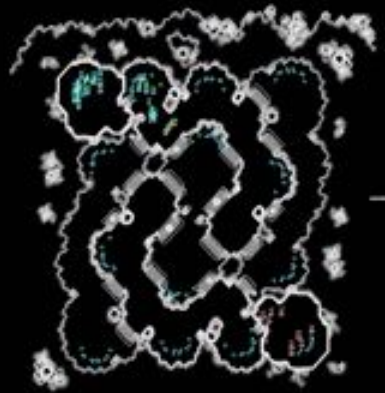

Neural Network Activations

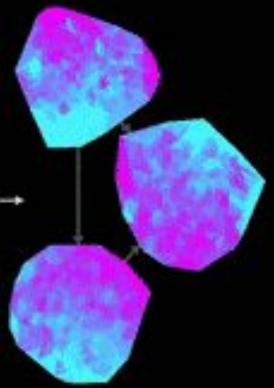

Considered Location

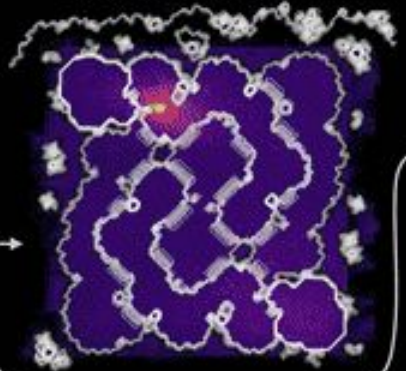

Outcome Prediction

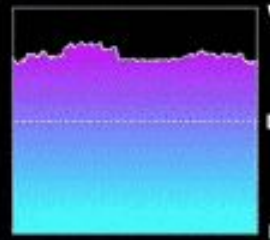

Considered Build/Train

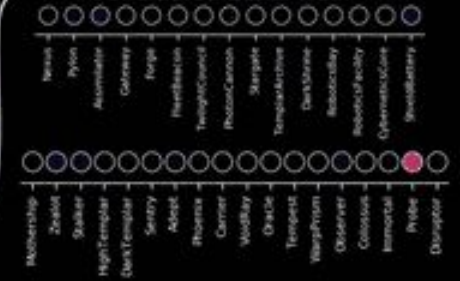

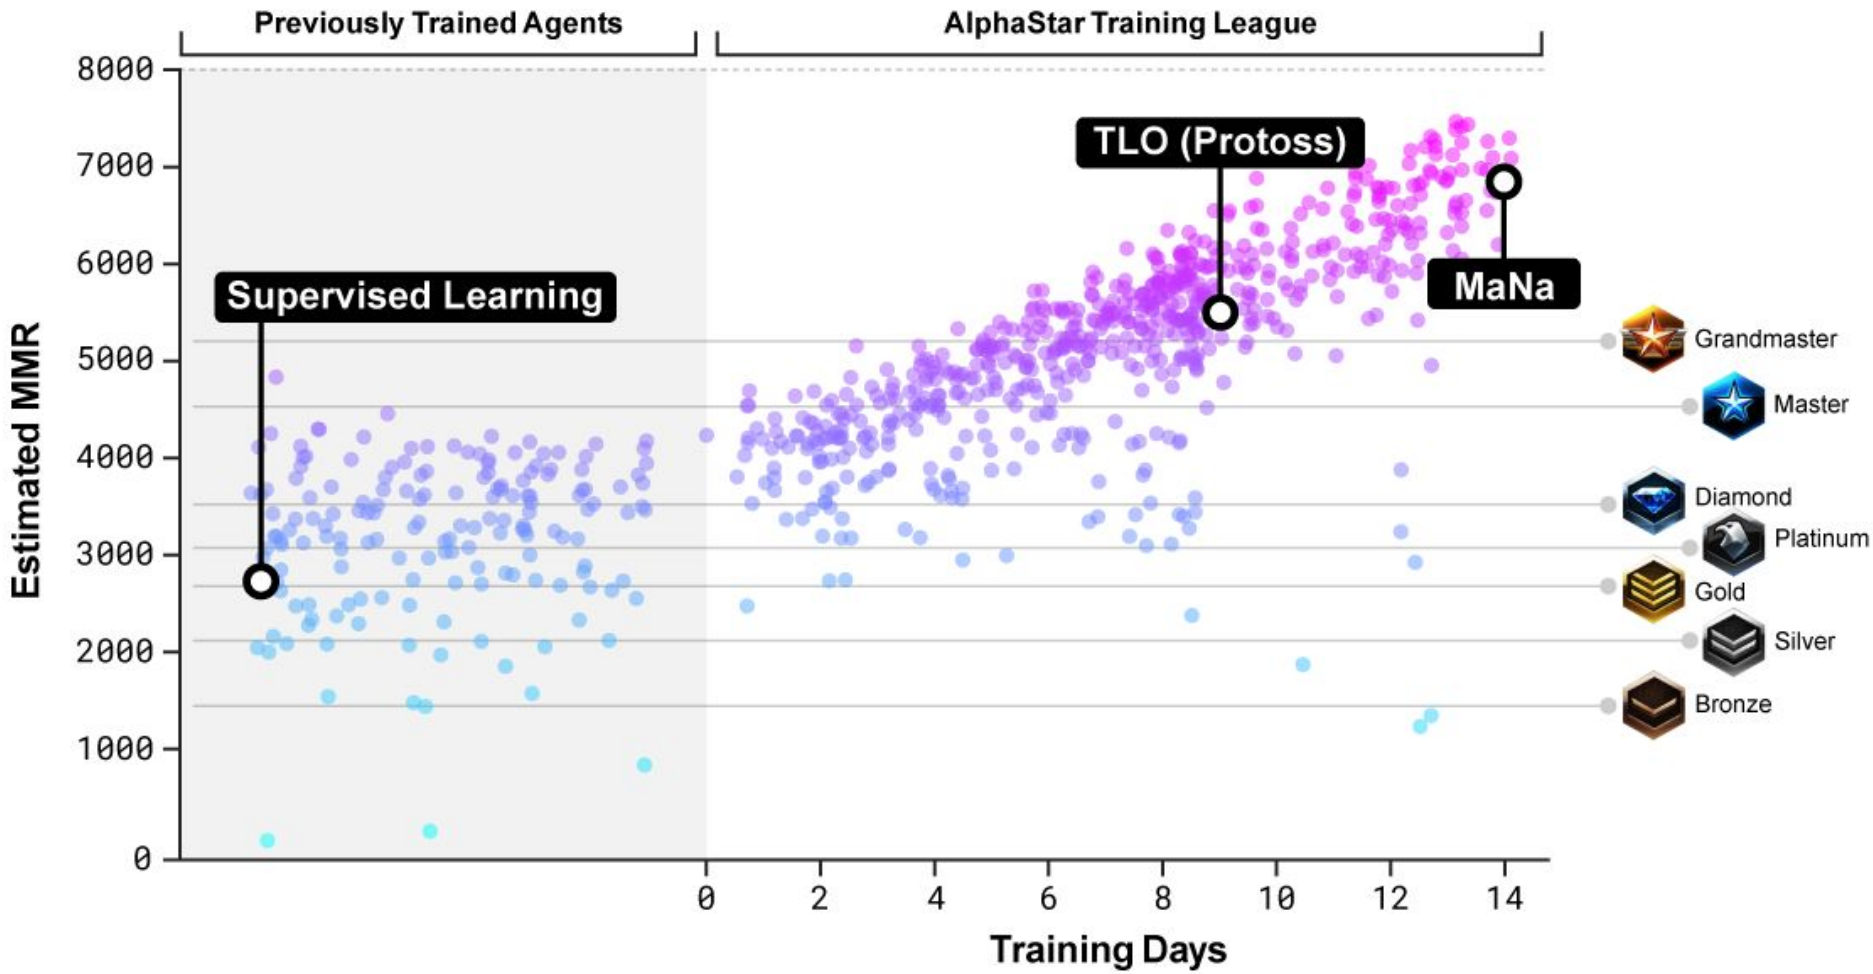
Previously Trained Agents
AlphaStar Training League
Supervised Learning
TLO (Protoss)
MaNa
Estimated MMR
8000
7000
6000
5000
4000
3000
2000
1000
0
0
2
4
6
8
10
12
14
Training Days
Grandmaster
Master
Diamond
Platinum
Gold
Silver
Bronze

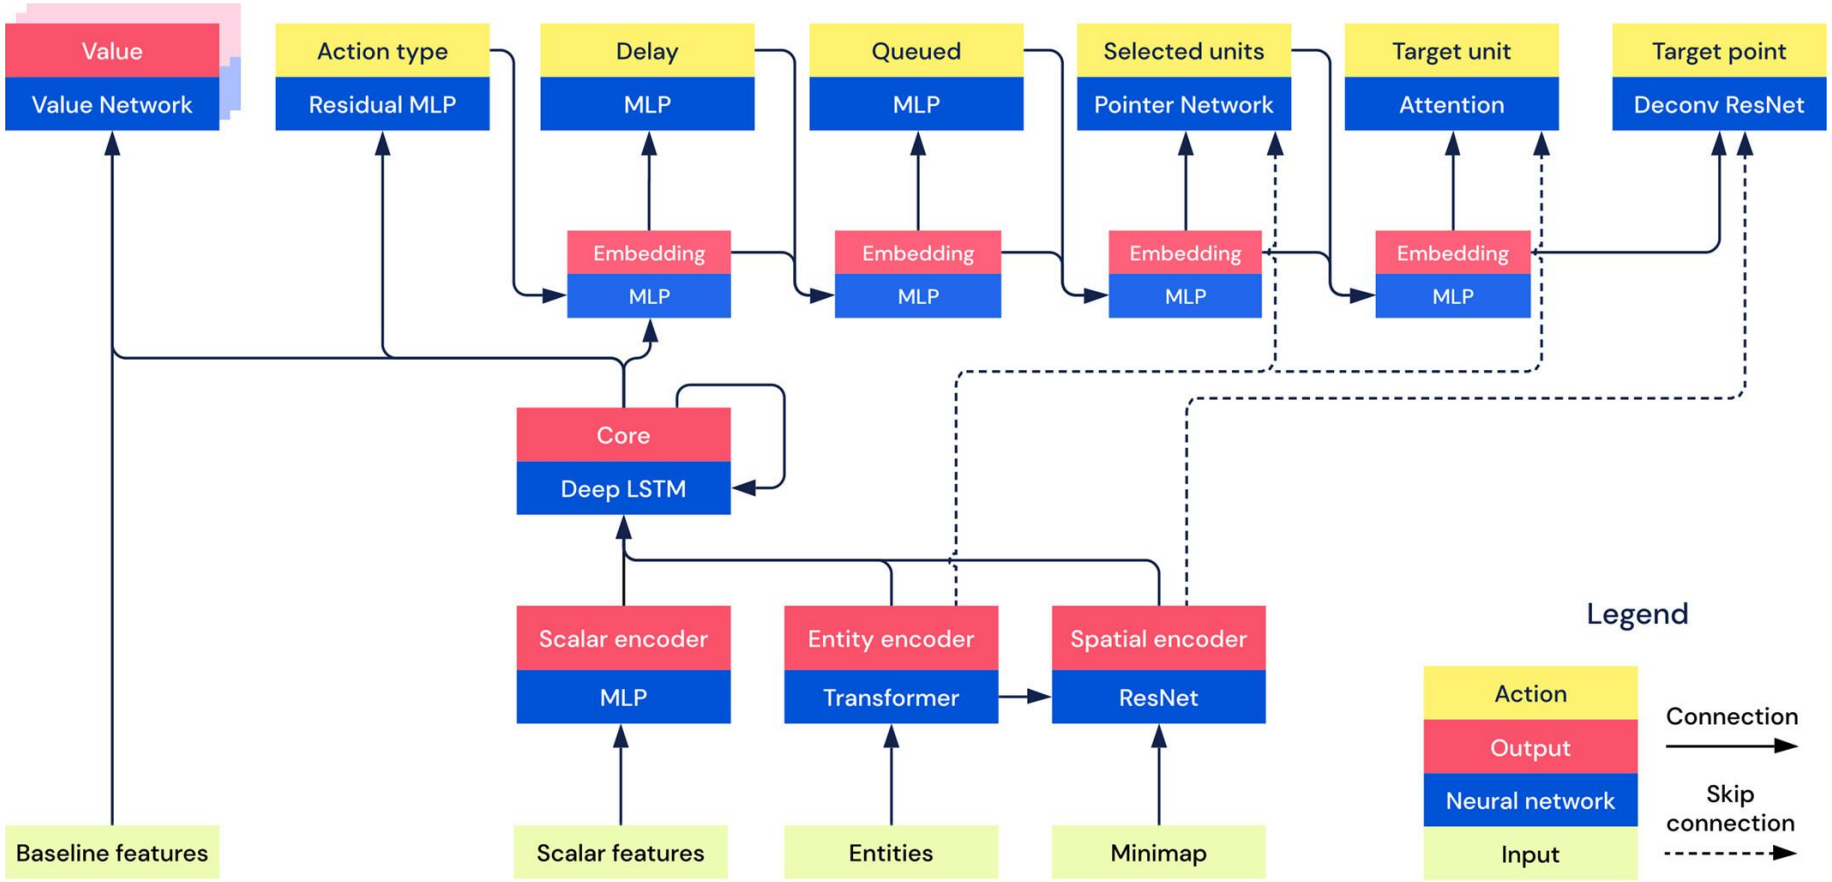
Value
Value Network
Action type
Residual MLP
Delay
MLP
Queued
MLP
Selected units
Pointer Network
Target unit
Attention
Target point
Deconv ResNet
Embedding
MLP
Embedding
MLP
Embedding
MLP
Embedding
MLP
Core
Deep LSTM
Scalar encoder
MLP
Entity encoder
Transformer
Spatial encoder
ResNet
Baseline features
Scalar features
Entities
Minimap
Legend
Action
Output
Neural network
Input
Connection
Skip connection

Raw Observations

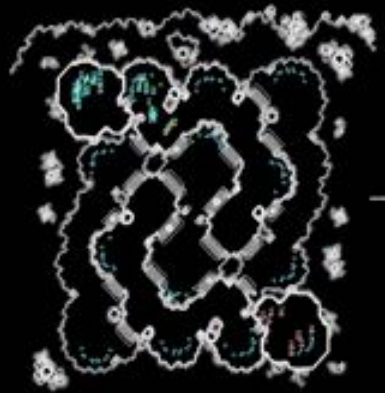

Neural Network Activations

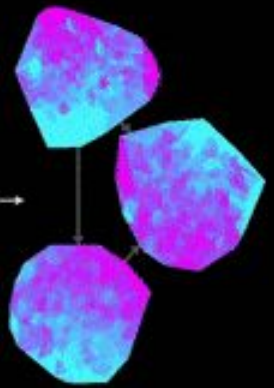

Considered Location

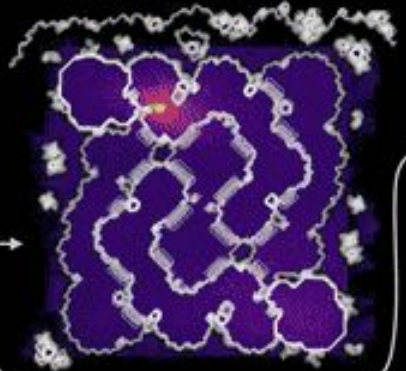

Outcome Prediction

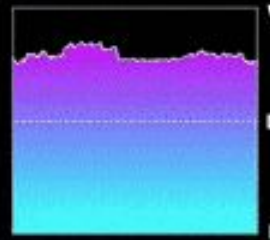

Considered Build/Train

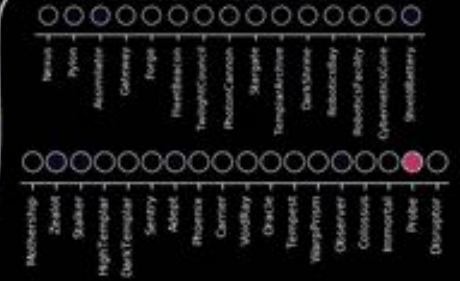

BLOG POST
RESEARCH

30 OCT 2019

AlphaStar: Grandmaster level in StarCraft II using multi-agent reinforcement learning

Article

Grandmaster level in StarCraft II using multi-agent reinforcement learning

https://doi.org/10.1038/s41586-019-1724-z

Received: 30 August 2019

Accepted: 10 October 2019

Published online: 30 October 2019

Oriol Vinyals[1,3]*, Igor Babuschkin[1,3], Wojciech M. Czarnecki[1,3], Michaël Mathieu[1,3], Andrew Dudzik[1,3], Junyoung Chung[1,3], David H. Choi[1,3], Richard Powell[1,3], Timo Ewalds[1,3], Petko Georgiev[1,3], Junhyuk Oh[1,3], Dan Horgan[1,3], Manuel Kroiss[1,3], Ivo Danihelka[1,3], Aja Huang[1,3], Laurent Sifre[1,3], Trevor Cai[1,3], John P. Agapiou[1,3], Max Jaderberg[1], Alexander S. Vezhnevets[1], Rémi Leblond[1], Tobias Pohlen[1], Valentin Dalibard[1], David Budden[1], Yury Sulsky[1], James Molloy[1], Tom L. Paine[1], Caglar Gulcehre[1], Ziyu Wang[1], Tobias Pfaff[1], Yuhuai Wu[1], Roman Ring[1], Dani Yogatama[1], Dario Wünsch[2], Katrina McKinney[1], Oliver Smith[1], Tom Schaul[1], Timothy Lillicrap[1], Koray Kavukcuoglu[1], Demis Hassabis[1], Chris Apps[1,3] & David Silver[1,3]*

Many real-world applications require artificial agents to compete and coordinate with other agents in complex environments. As a stepping stone to this goal, the domain of StarCraft has emerged as an important challenge for artificial intelligence research, owing to its iconic and enduring status among the most difficult professional esports and its relevance to the real world in terms of its raw complexity and multi-agent challenges. Over the course of a decade and numerous competitions[1–3], the strongest agents have simplified important aspects of the game, utilized superhuman capabilities, or employed hand-crafted sub-systems[4]. Despite these advantages, no previous agent has come close to matching the overall skill of

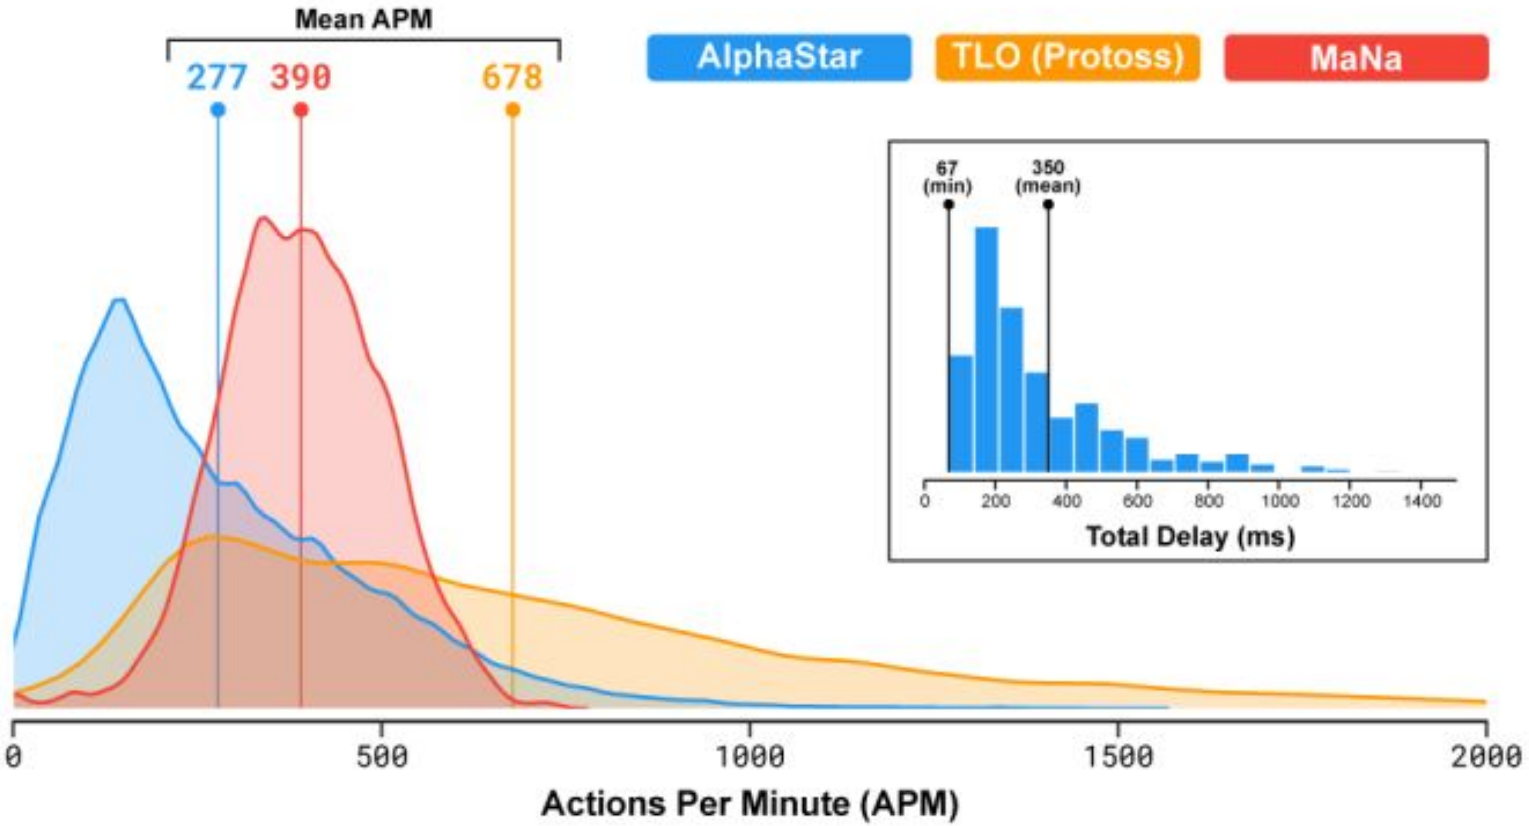

CLARIFICATION (29/01/19): TLO'S APM APPEARS HIGHER THAN BOTH ALPHASTAR AND MANA BECAUSE OF HIS USE OF RAPID-FIRE HOT-KEYS AND USE OF THE "REMOVE AND ADD TO CONTROL GROUP" KEY BINDINGS. ALSO NOTE THAT ALPHASTAR'S EFFECTIVE APM BURSTS ARE SOMETIMES HIGHER THAN BOTH PLAYERS.

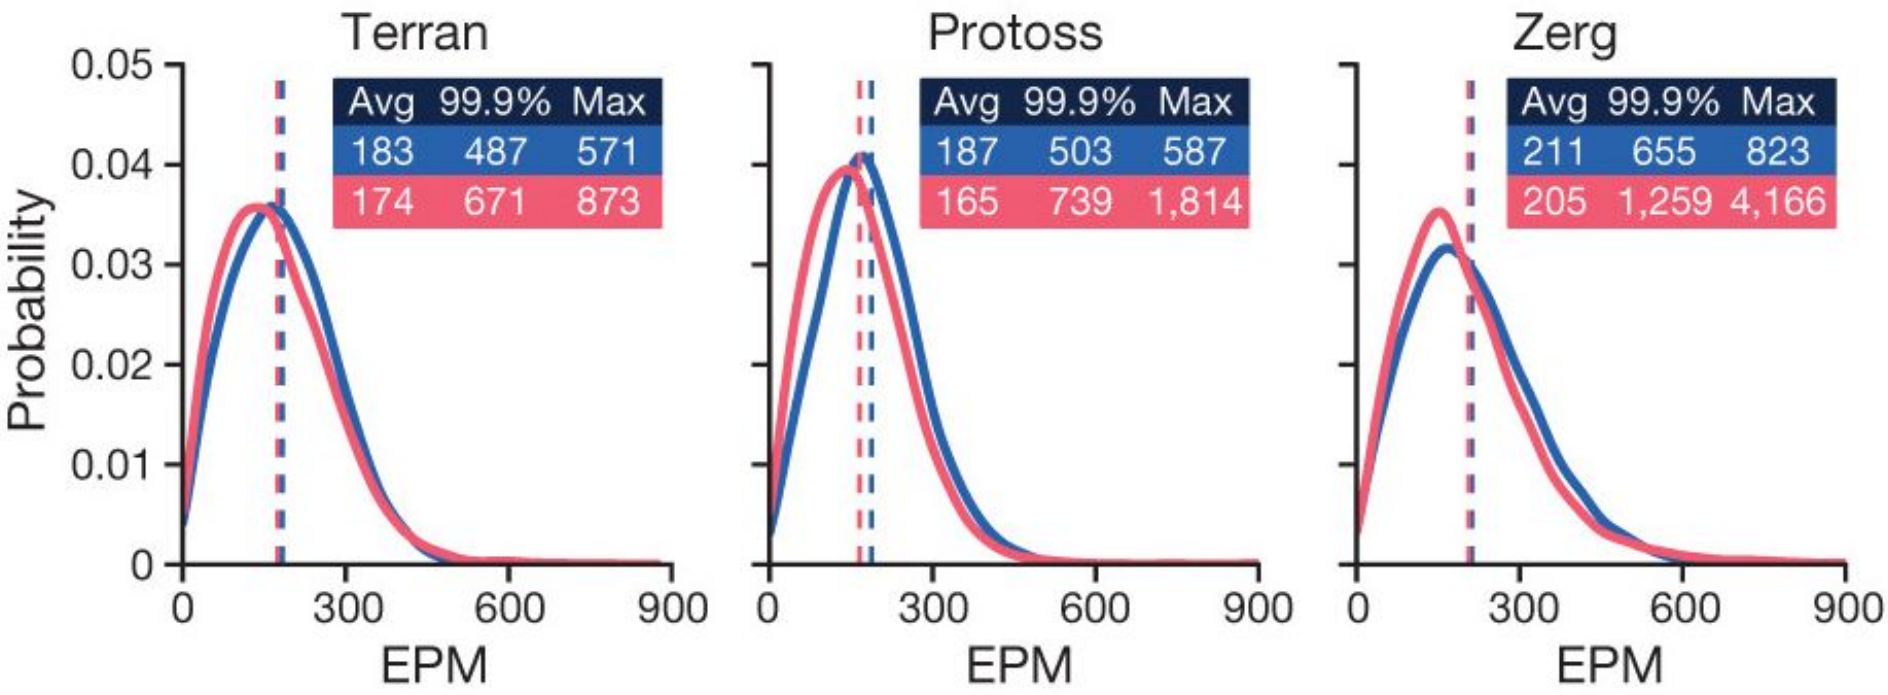
Terran
Avg 99.9% Max
183 487 571
174 671 873
Protoss
Avg 99.9% Max
187 503 587
165 739 1,814
Zerg
Avg 99.9% Max
211 655 823
205 1,259 4,166
Probability
0.05
0.04
0.03
0.02
0.01
0
0 300 600 900
EPM


PRODUCTION
Workers: 4/16
Workers: 0/3
5:57
Dusk Towers
BEST OF 3
0 <FCEO>PandaBear... 140 204 59/68 35 24
0 <Grav>EJK 351 300 66/61 38 26
Nydus Worm
Current Status


Render of Agent's view
AlphaStar


MaNa

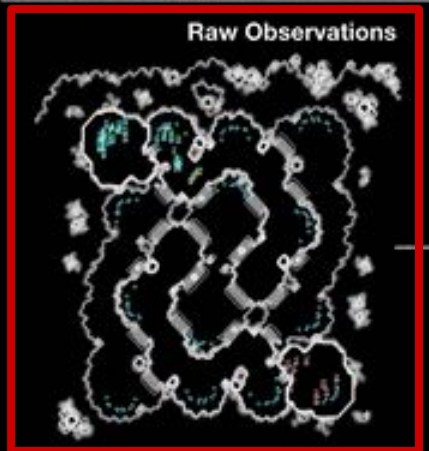
Raw Observations

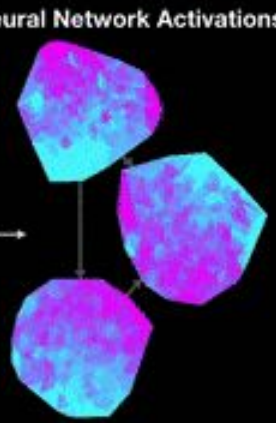
Neural Network Activations

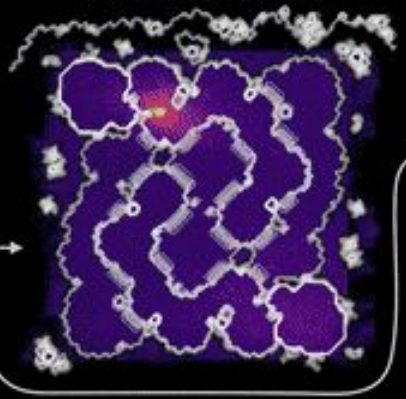
Considered Location

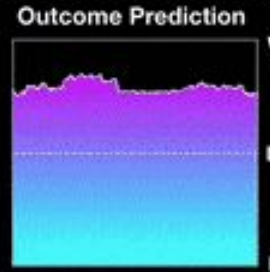
Outcome Prediction
Win
Draw
Lose

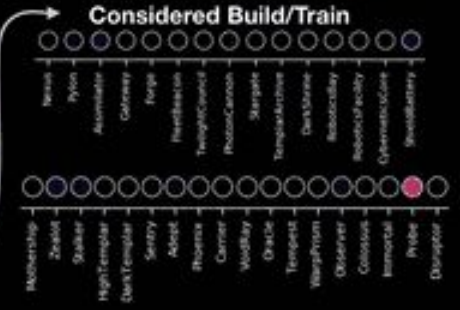
Considered Build/Train

JANUARY 24

DeepMind's AlphaStar AI wins 10-1 against professional StarCraft II players

Abner Li - Jan. 24th 2019 12:50 pm PT @technacity

DeepMind AI AlphaStar goes 10-1 against top 'StarCraft II' pros

The AI beat 'StarCraft' pros TLO and MaNa thanks to more than 200 years worth of game knowledge.

AJ Dellinger, @ajdell
01.24.19 in Robots

24
Comments

2743
Shares

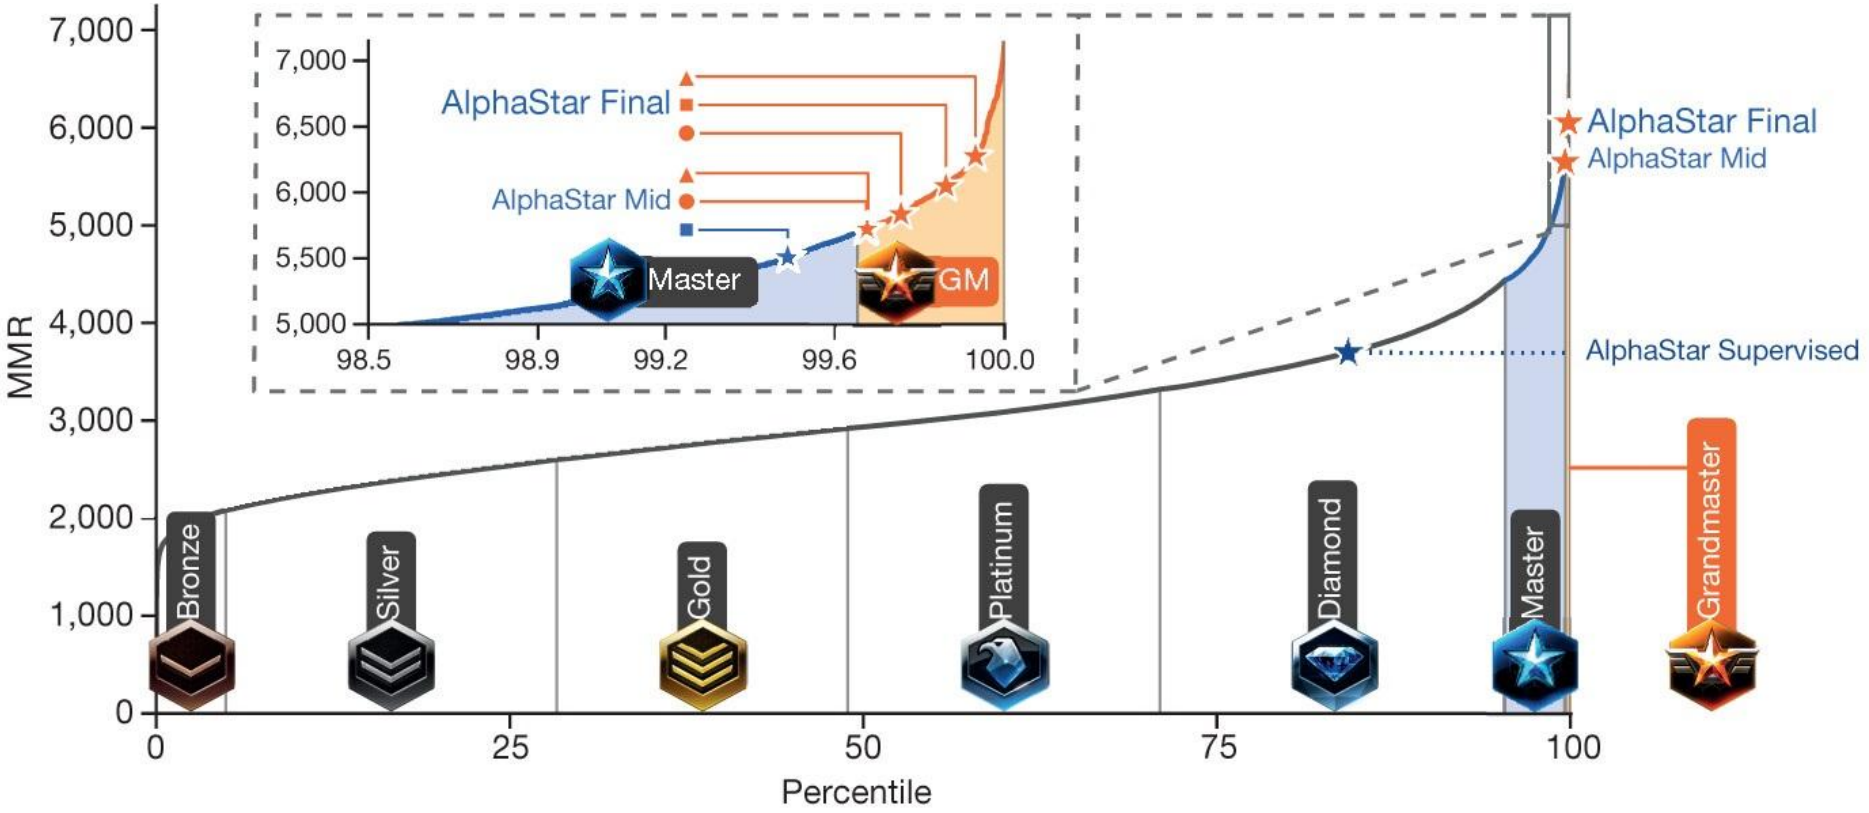
7,000
6,000
5,000
4,000
3,000
2,000
1,000
0
MMR
0
25
50
75
100
Percentile
Bronze
Silver
Gold
Platinum
Diamond
Master
Grandmaster
AlphaStar Final
AlphaStar Mid
AlphaStar Supervised
7,000
6,500
6,000
5,500
5,000
98.5
98.9
99.2
99.6
100.0
AlphaStar Final
AlphaStar Mid
Master
GM

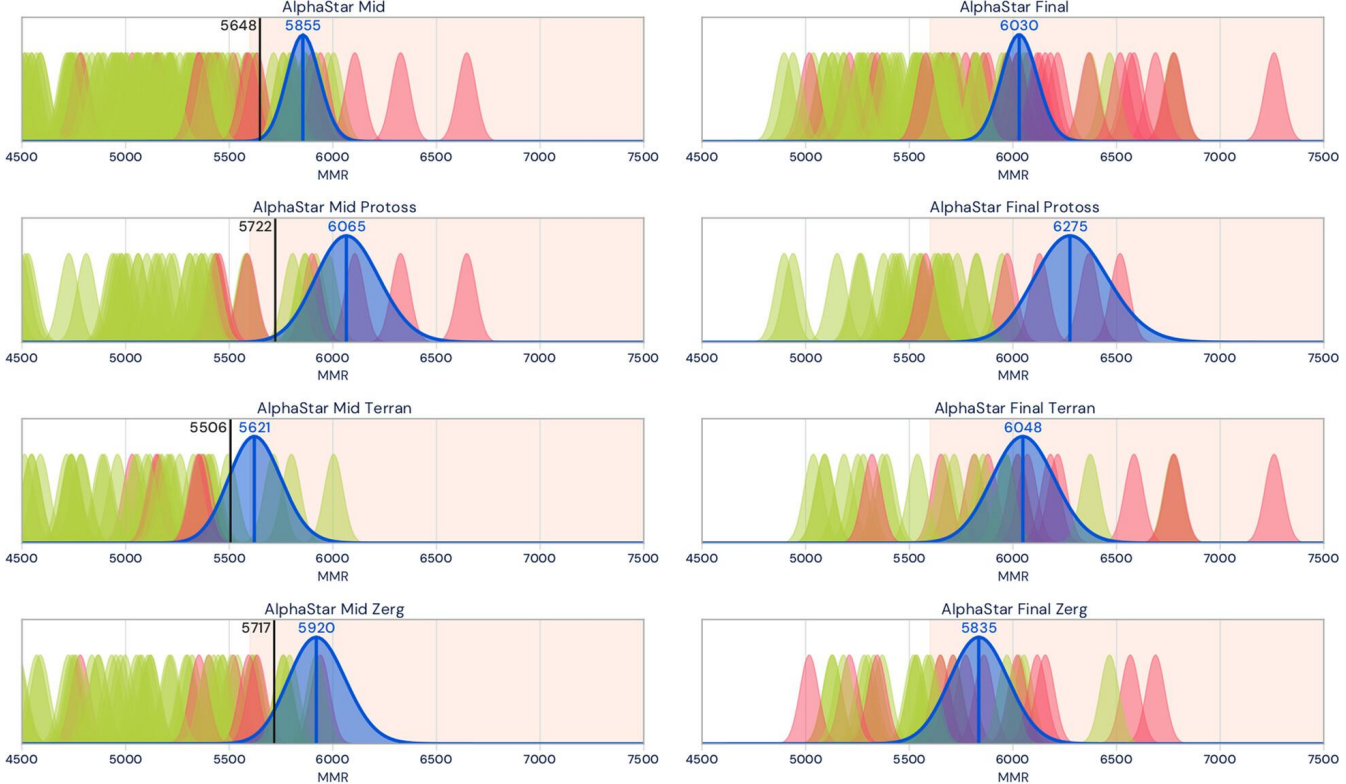
AlphaStar Mid
5648
5855
AlphaStar Final
6030
AlphaStar Mid Protoss
5722
6065
AlphaStar Final Protoss
6275
AlphaStar Mid Terran
5506
5621
AlphaStar Final Terran
6048
AlphaStar Mid Zerg
5717
5920
AlphaStar Final Zerg
5835
4500
5000
5500
6000
6500
7000
7500
MMR

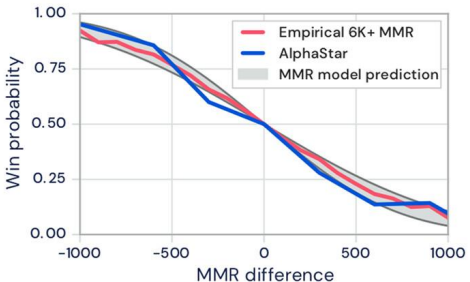
Empirical 6K+ MMR
AlphaStar
MMR model prediction
Win probability
1.00
0.75
0.50
0.25
0.00
-1000
-500
0
500
1000
MMR difference

DeepMind | StarCraft II

DEMONSTRATION

GRZEGORZ 'MANA' KOMINCZ

“Foi animador ver o agente desenvolver suas próprias estratégias de maneira diferente dos jogadores humanos [...]. Os limites nas ações que o agente pode executar e a restrição na visão da câmera agora tornam as partidas convincentes - embora, como um profissional, eu ainda possa visualizar algumas das fraquezas do sistema” [6]

BLIZZARD ENTERTAINMENT


DeepMind
StarCraft II
DEMONSTRATION
BLIZZARD
ENTERTAINMENT

NEWS

09 AUG 2017

SHARE

AUTHOR

OV Oriol Vinyals
SG Stephen Gaffney
TE Timo Ewalds

DeepMind and Blizzard open StarCraft II as an AI research environment

DeepMind's scientific mission is to push the boundaries of AI by developing systems that can learn to solve complex problems. To do this, we design agents and test their ability in a wide range of environments from the purpose-built DeepMind Lab to established games, such as Atari and Go.

Testing our agents in games that are not specifically designed for AI research, and where humans play well, is crucial to benchmark agent performance. That is why we, along with our partner Blizzard Entertainment, are excited to announce the release of SC2LE, a set of tools that we hope will accelerate AI research in the real-time

StarCraft II: A New Challenge for Reinforcement Learning

Oriol Vinyals **Timo Ewalds** **Sergey Bartunov** **Petko Georgiev**
Alexander Sasha Vezhnevets **Michelle Yeo** **Alireza Makhzani** **Heinrich Küttler**
John Agapiou **Julian Schrittwieser** **Stephen Gaffney** **Stig Petersen**
Karen Simonyan **Tom Schaul** **Hado van Hasselt** **David Silver** **Timothy Lillicrap**
DeepMind

Kevin Calderone **Paul Keet** **Anthony Brunasso** **David Lawrence**
Anders Ekermo **Jacob Repp** **Rodney Tsing**
Blizzard

Abstract

This paper introduces *SC2LE* (StarCraft II Learning Environment), a reinforcement learning environment based on the StarCraft II game. This domain poses a new grand challenge for reinforcement learning, representing a more challenging class of problems than considered in most prior work. It is a multi-agent problem with multiple players interacting; there is imperfect information due to a partially observed map; it has a large action space involving the selection and control of hundreds of units; it has a large state space that must be observed solely from raw input feature planes; and it has delayed credit assignment requiring long-term strategies over thousands of steps. We describe the observation, action, and reward

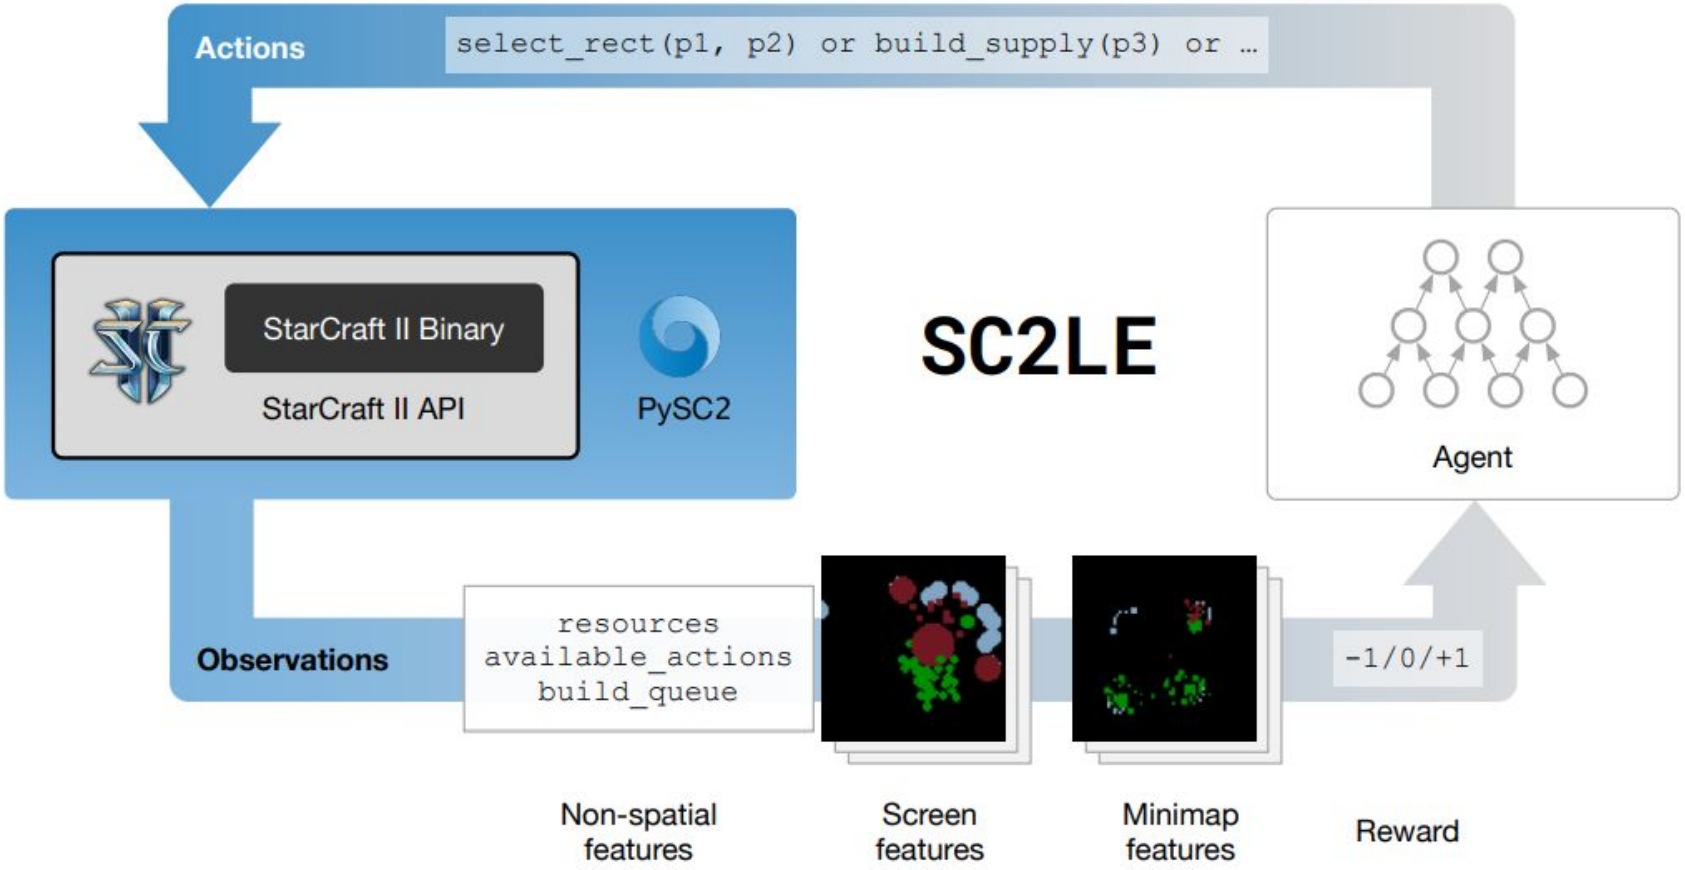
Actions
select_rect(p1, p2) or build_supply(p3) or …
StarCraft II Binary
StarCraft II API
PySC2
SC2LE
Agent
Observations
resources
available_actions
build_queue
-1/0/+1
Non-spatial features
Screen features
Minimap features
Reward

README.md	Simplify from source installation example.	9 months ago
setup.py	Bump the version to 3.0.0 for a public release.	2 months ago

README.md

PySC2 - StarCraft II Learning Environment

PySC2 is DeepMind's Python component of the StarCraft II Learning Environment (SC2LE). It exposes Blizzard Entertainment's StarCraft II Machine Learning API as a Python RL Environment. This is a collaboration between DeepMind and Blizzard to develop StarCraft II into a rich environment for RL research. PySC2 provides an interface for RL agents to interact with StarCraft 2, getting observations and sending actions.

We have published an accompanying blogpost and paper, which outlines our motivation for using StarCraft II for DeepRL research, and some initial research results using the environment.

About

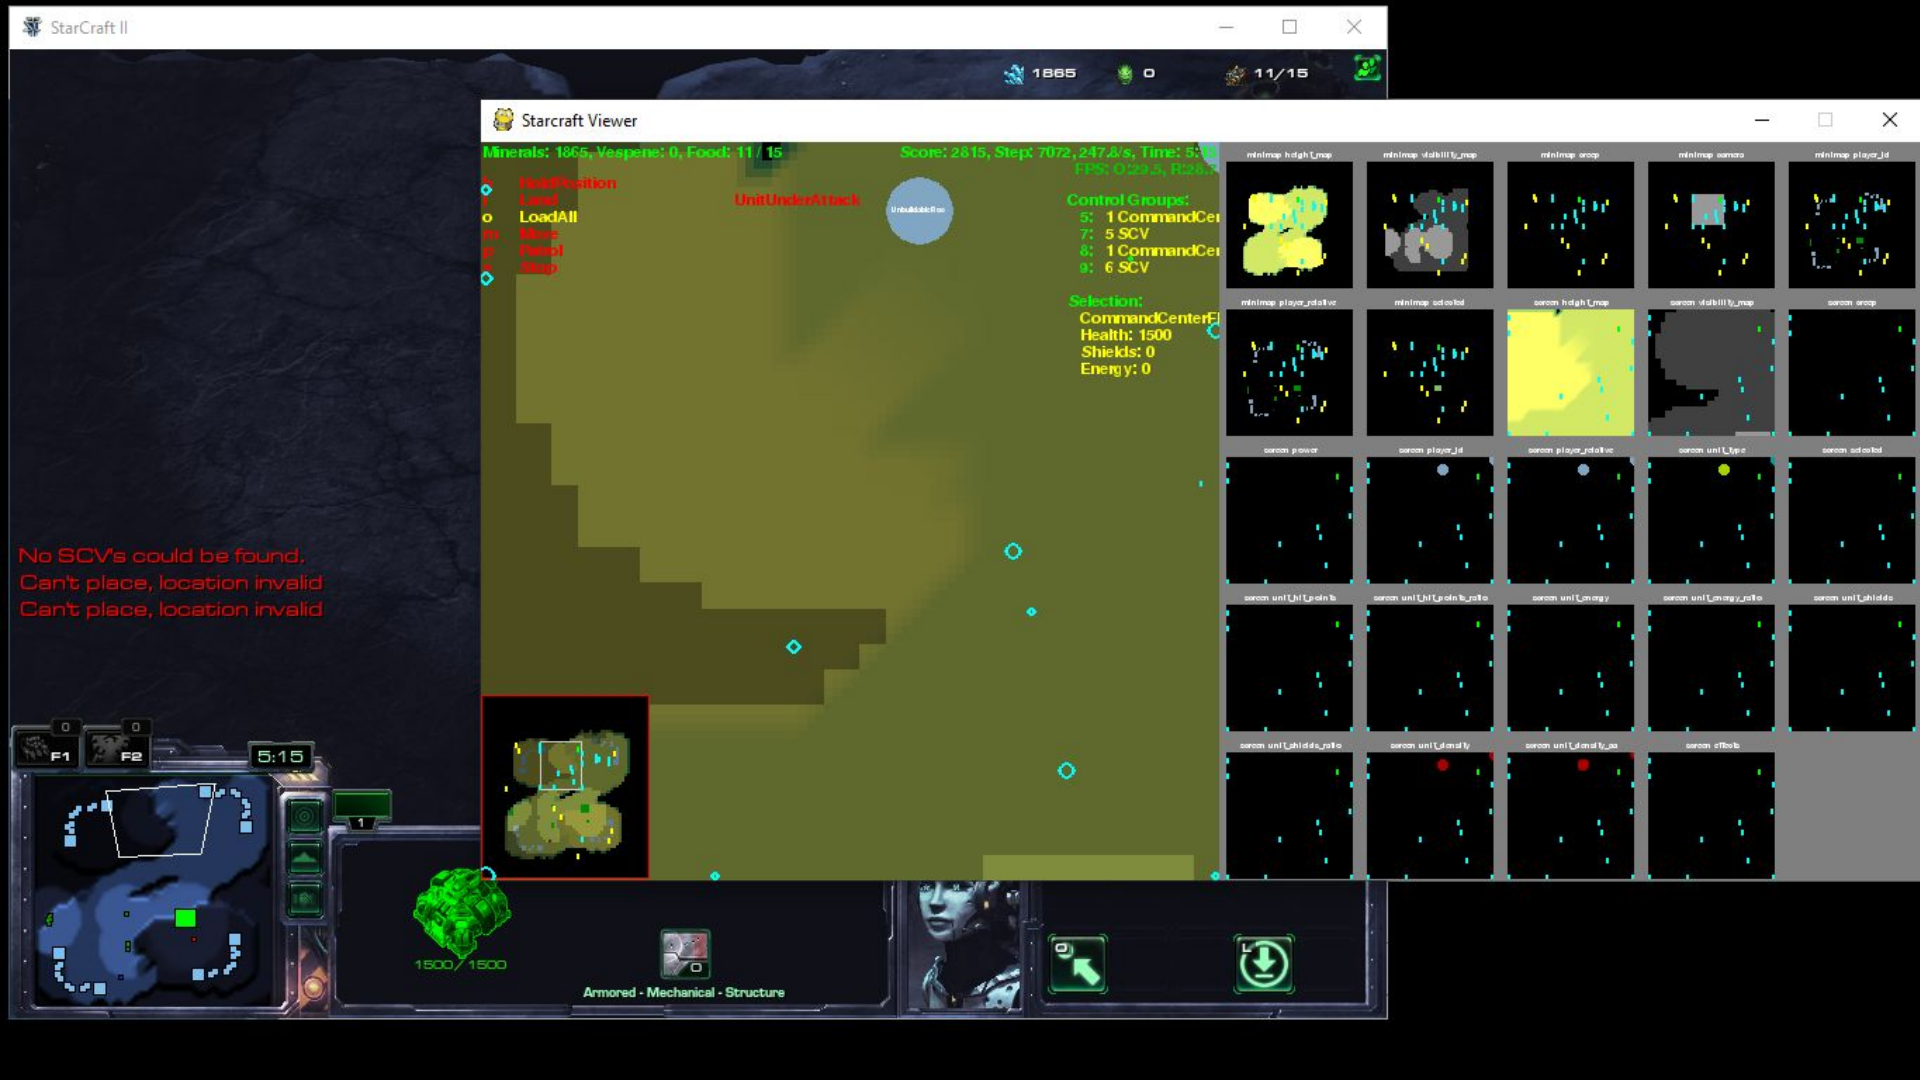
StarCraft II
1865
0
11/15
Starcraft Viewer
Minerals: 1865, Vespene: 0, Food: 11 / 15
Score: 2815, Step: 7072, 247.8/s, Time: 5:
FPS: O:29.5, R:28.
HoldPosition
Land
LoadAll
Move
Patrol
Stop
UnitUnderAttack
Control Groups:
5: 1 CommandCe
7: 5 SCV
8: 1 CommandCe
9: 6 SCV
Selection:
CommandCenterF
Health: 1500
Shields: 0
Energy: 0
minimap height_map
minimap visibility_map
minimap creep
minimap camera
minimap player_id
minimap player_relative
minimap selected
screen height_map
screen visibility_map
screen creep
screen power
screen player_id
screen player_relative
screen unit_type
screen selected
screen unit_hit_points
screen unit_hit_points_ratio
screen unit_energy
screen unit_energy_ratio
screen unit_shields
screen unit_shields_ratio
screen unit_density
screen unit_density_ra
screen effects
No SCV's could be found.
Can't place, location invalid
Can't place, location invalid
F1
F2
5:15
1500 / 1500
Armored - Mechanical - Structure

Human Actions

Agent Actions

Available Actions


Move to Beacon


Collect Minerals and Gas


Collect Mineral Shards


Build Marines

Untrained Agent

Trained Supervised Agent

Copyright

Links recomendados

DeepMind blog posts

https://deepmind.com/blog/announcements/deepmind-and-blizzard-open-starcraft-ii-ai-research-environment
https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii
https://deepmind.com/blog/article/AlphaStar-Grandmaster-level-in-StarCraft-II-using-multi-agent-reinforcement-learning

Canal DeepMind no YouTube

https://www.youtube.com/channel/UCP7jMXSY2xbc3KCAE0MHQ-A

Análises AlphaStar vs Serral na Blizzcon 2019, por Artosis

https://youtu.be/OxseexGkv_Q

AlphaStar líderes de pesquisa - “Ask Me Anything” no Reddit

https://www.reddit.com/r/MachineLearning/comments/ajgzoc/we_are_oriol_vinyals_and_david_silver_from/

Análise Crítica das partidas contra TLO e MaNa, por Aleksi Pietikäinen

https://blog.usejournal.com/an-analysis-on-how-deepminds-starcraft-2-ai-s-superhuman-speed-could-be-a-band-aid-fix-for-the-1702fb8344d6

Links das fontes

Papers

Sutton and Barto. Reinforcement Learning: An Introduction. 2018.
http://incompleteideas.net/book/the-book.html

Vinyals et al. StarCraft II: A New Challenge for Reinforcement Learning. 2017.
https://arxiv.org/abs/1708.04782

Vinyals et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. Nature. 2019.
https://www.nature.com/articles/s41586-019-1724-z.epdf?author_access_token=lZH3nqPYtWJXfDA10W0CNNRgN0jAjWel9jnR3ZoTv0PSZcPzJFGNAZhOlk4deBCKzKm70KfinloafEF1bCCXL6IIHHgKaDkaTkBcTEv7aT-wqDoG1VeO9-wO3GEoAMF9bAOt7mJ0RWQnRVMbyfgH9A%3D%3D

Mnih et al. Playing Atari with Deep Reinforcement Learning. 2013.
https://arxiv.org/abs/1312.5602

Mnih et al. Human-level control through deep reinforcement learning. Nature. 2015.
https://www.nature.com/articles/nature14236

Links das fontes

Referências

[1] StarCraft II 3 milhões de cópias vendidas
https://www.eurogamer.net/articles/2010-09-01-starcraft-ii-sells-3-million-in-a-month

[2] TLO antes das partidas
https://youtu.be/UuhECwm31dM?t=87

[3] TLO após as partidas
https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii

[4] MaNa antes das partidas
https://youtu.be/UuhECwm31dM?t=215

[5] MaNa após as partidas
https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii

[6] MaNa sobre nova versão AlphaStar
https://deepmind.com/blog/article/AlphaStar-Grandmaster-level-in-StarCraft-II-using-multi-agent-reinforcement-learning

Links das fontes

Sites

Github PySC2
https://github.com/deepmind/pysc2

AlphaStar 10 - 1 TLO/MaNa
https://www.engadget.com/2019/01/24/deepmind-ai-starcraft-ii-demonstration-tlo-mana/
https://tl.net/forum/starcraft-2/541114-alphastar-ai-goes-10-1-against-human-pros-in-demonstration
https://9to5google.com/2019/01/24/deepmind-alphastar-wins-starcraft-ii/

Notícias AlphaStar Grande Mestre
https://www.bbc.com/news/technology-50212841
https://www.sciencealert.com/starcraft-ii-has-a-new-grandmaster-and-it-s-not-human?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%3A+sciencealert-latestnews+%28ScienceAlert-Latest%29
https://www.theguardian.com/technology/2019/oct/30/ai-becomes-grandmaster-in-fiendishly-complex-starcraft-ii
https://olhardigital.com.br/games-e-consoles/noticia/inteligencia-artificial-do-google-vence-99-8-das-partidas-de-starcraft-ii/92316
https://canaltech.com.br/inteligencia-artificial/starcraft-inteligencia-artificial-deixa-998-dos-jogadores-no-chinelo-153967/
https://www.tecmundo.com.br/software/147376-ia-deepmind-vencer-99-8-jogadores-starcraft-2.htm

Links das fontes

Imagens/animações

StarCraft II gif
https://media.giphy.com/media/FbUbolChwFN6M/source.gif

Complexidade dos jogos
https://twitter.com/liv_boeree/status/1045025689710268421

Processo de Decisão de Markov (MDP)
https://en.wikipedia.org/wiki/Markov_decision_process

Mario MDP
https://paulovasconcellos.com.br/explicando-deep-reinforcement-learning-com-super-mario-ao-inv%C3%A9s-de-matem%C3%A1tica-4c77392cc733

Interação Ambiente-Agente Rede Neural
http://people.csail.mit.edu/hongzi/content/publications/DeepRM-HotNets16.pdf

Links das fontes

Vídeos

AphaStar vs TLO/MaNa
https://youtu.be/cUTMhmVh1qs

PySC2 ambiente
https://youtu.be/-fKUyT14G-8

PySC2 minigames
https://youtu.be/6L448yg0Sm0

Trained vs. untrained agent
https://youtu.be/WEOzide5XFc

DQN Breakout
https://youtu.be/TmPfTpjtdgg

Muito Obrigado!

Paulo Bruno de Sousa Serafim

paulo_serafim@atlantico.com.br
paulobruno@alu.ufc.br

Atlântico

Porto Alegre
2019

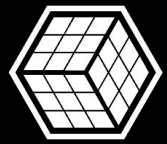

THE DEVELOPER'S CONFERENCE