



AlphaStar: Mastering the Real-Time Strategy Game Starcraft II

Johannes Daub

AI for Games – 11.07.19



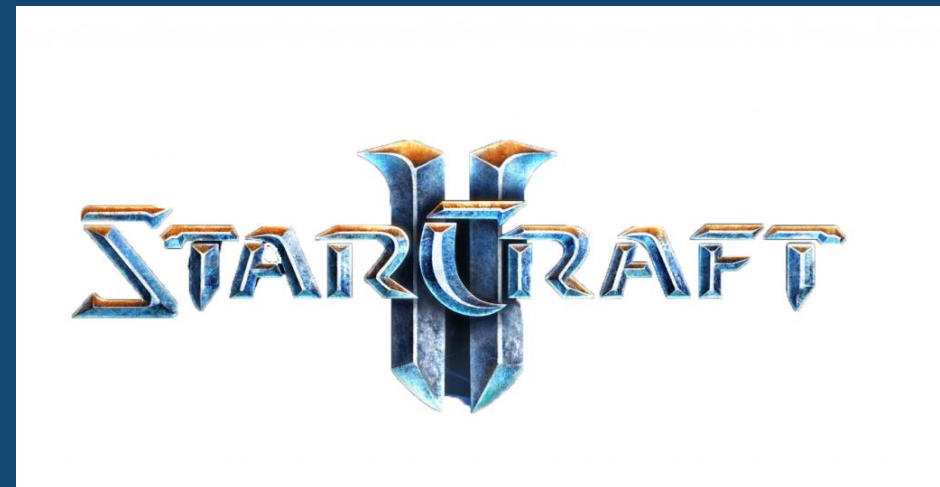
Content

11.07.2019

- Introduction
- Part I – 2017: The Beginning
 - Framework
 - Mini-Games
 - Evaluation
- Part II – 2019: The Mastery
 - AlphaStar

Starcraft II

- Real-Time Strategy
- Made by Blizzard Entertainment
- Sci-Fi Theme
- 3 Races with completely different playstyles
- Competitive Scene



[1]

Protoss



[2]

Zerg & Terran



[3]

Google Deepmind Team



[4]

Oriol Vinyals



[5]

- Part of Google Brain before
- His research is used in Google Translate, Text-To-Speech and Speech recognition
- Cited over 43000 times

David Silver



[6]

- Professor of Computer Science of University College London
- Lead researcher of AlphaGo/AlphaZero
- Cited over 29000 times



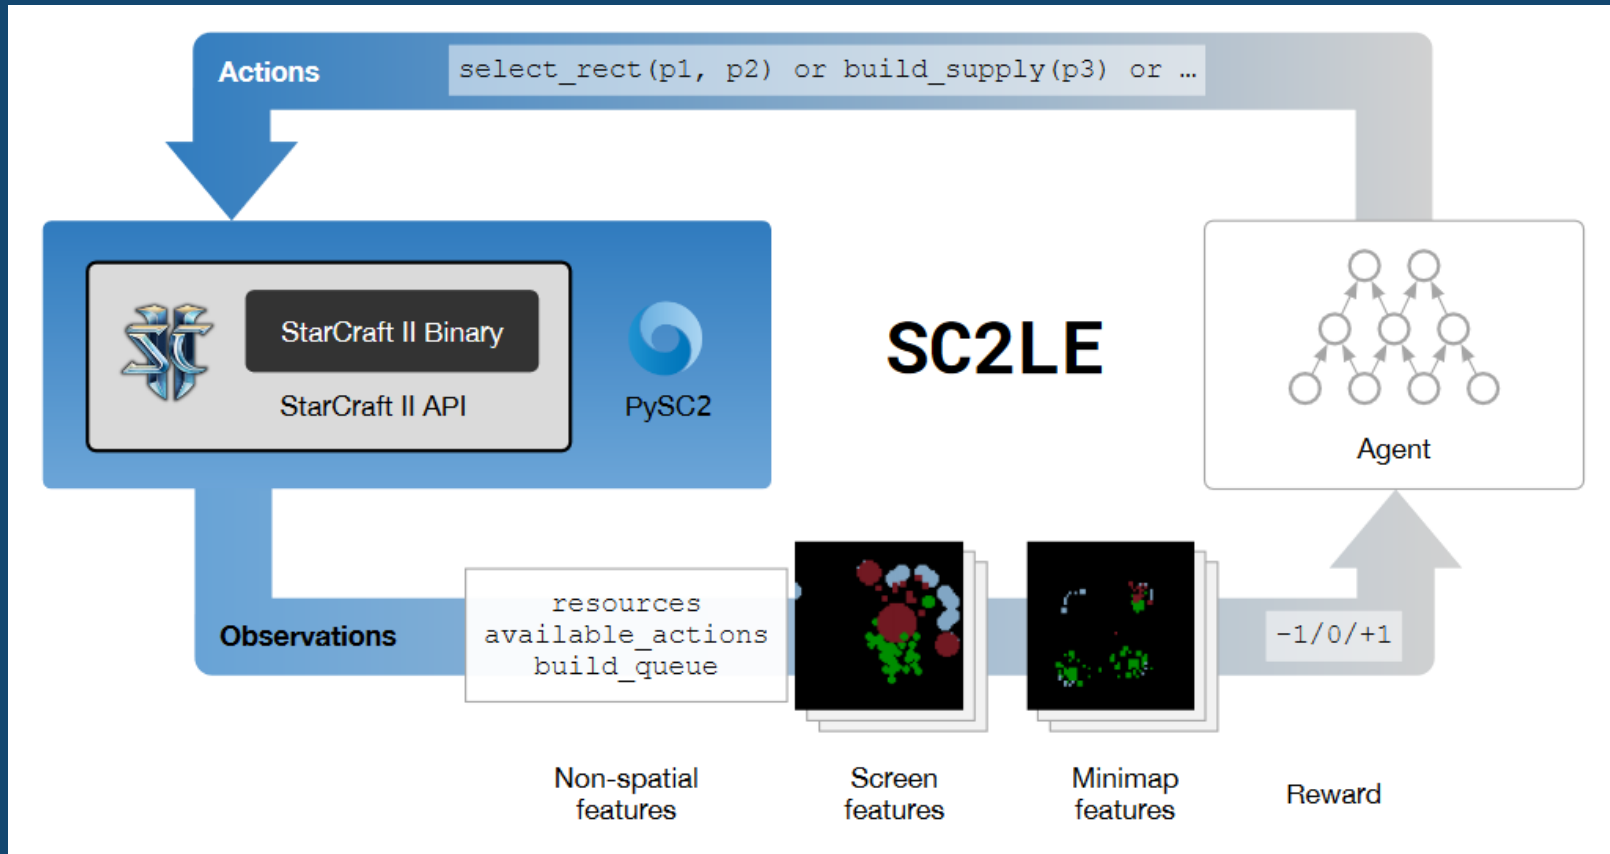
THE BEGINNING

Part I

Why Starcraft?

- **Real time:** Continuous Action required
- **Imperfect information:** Only part of the game state visible
- **Long term planning:** Early actions may payoff later
- **Large action space**
- **Game theory:** There is no single superior strategy (rock-paper-scissors)

SC2LE – Starcraft 2 Learning Environment [7]



[7]

Observations I

- Use feature layers instead of 3D image
 - Main map
 - Minimap
 - Interface



[8]

Observations II

Starcraft Viewer
 Minerals: 4250, Vespene: 789, Food: 108 / 110; Score: 16564, Frame: 14377, FPS: G:44.4, R:26.0

- a - AcquireMove
- a - Attack
- cc - Colossus
- cd - WarpInDisruptor
- ce - Probe
- ch - WarpInAdept
- ei - Immortal
- es - Oracle
- en - Archon
- cr - WarpPrism
- cs - Stalker
- ct - HighTemplar
- cv - VoidRay
- cx - Phoenix
- cz - Zealot
- e - PhasingMode
- e - VoidRaySwarmDamageBoost
- f - ForceField
- f - MothershipCoreWeapon
- f - AdeptPhaseShift
- g - GuardianShield

ssRecall

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_hill_points	screen unit_energy	screen unit_shields	screen unit_density	screen unit_density_as

[7]


Actions

Human Actions	IDLE	Left_Click_Hold (p1)	Press B + S	IDLE
		Release (p2)	Left_Click (p3)	
Agent Actions	no_op	select_rect(p1, p2)	build_supply(p3)	no_op
Available Actions	<p>Base action Point Point</p> <p>no_op </p> <p>rectangle select </p>	<p>Base action Point Point</p> <p>no_op </p> <p>rectangle select </p>	<p>Base action Point Point</p> <p>no_op </p> <p>rectangle select </p> <p>Build supply </p>	<p>Base action Point Point</p> <p>no_op </p> <p>rectangle select </p> <p>Build supply </p>

[7]

Mini Games

- **MoveToBeacon:** Get score for reaching a beacon with a unit (+1)
- **FindAndDefeatZerglings:** Move units and defeat enemies (+2)
- **BuildMarines:** Build workers, collect resources, build Supply Depots, build Barracks, and then train marines. (+1)

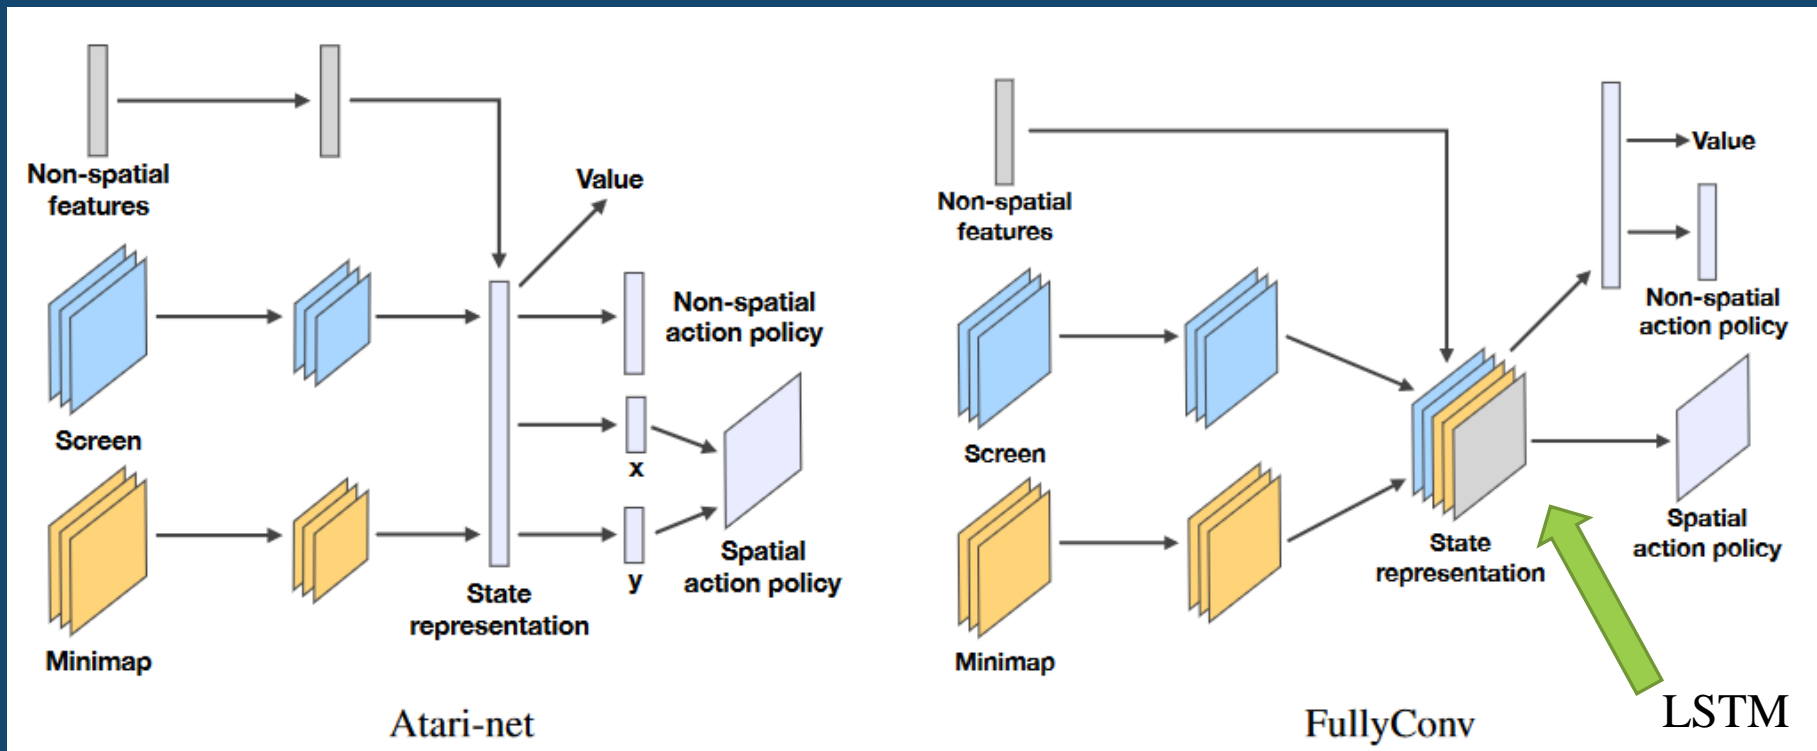


[9]

Baseline Agents


- **Atari-net Agent:** Also used for Atari Benchmark. CNN + FC
- **FullyConv Agent:** Similar architecture, but preserving spatial structure
- **FullyConv LSTM Agent:** Add a LSTM for memory

Baseline Agents



[7]

Performance on Mini Games



[7]

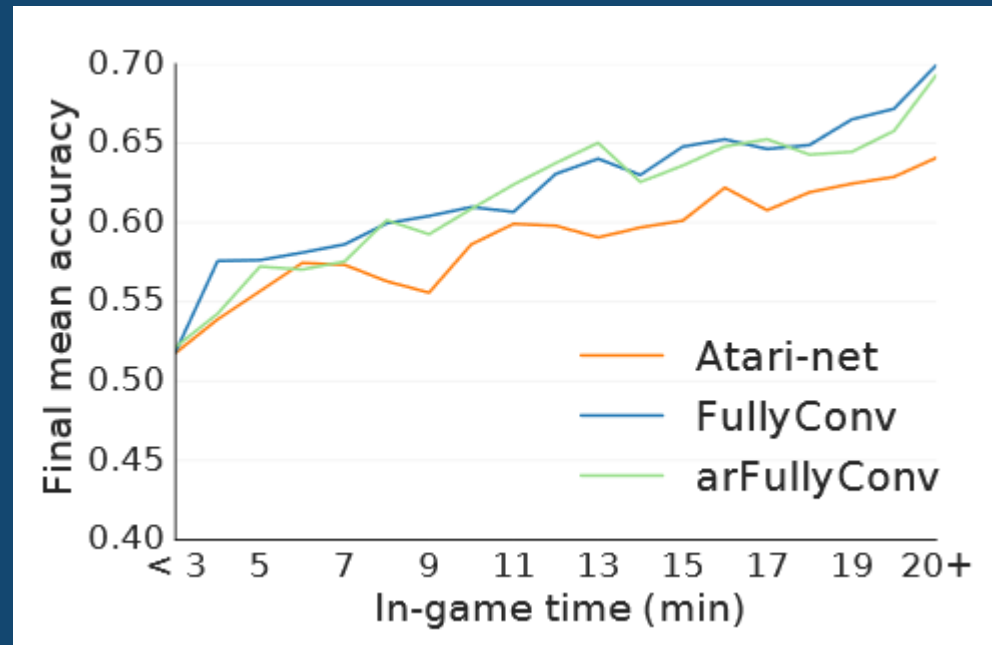
Performance on Mini Games

AGENT	METRIC	MOVE TO BEACON	COLLECT MINERAL SHARDS	FIND AND DEFEAT ZERGLINGS	DEFEAT ROACHES	DEFEAT ZERGLINGS AND BANELINGS	COLLECT MINERALS AND GAS	BUILD MARINES
RANDOM POLICY	MEAN	1	17	4	1	23	12	< 1
	MAX	6	35	19	46	118	750	5
RANDOM SEARCH	MEAN	25	32	21	51	55	2318	8
	MAX	29	57	33	241	159	3940	46
DEEPMIND HUMAN PLAYER	MEAN	26	133	46	41	729	6880	138
	MAX	28	142	49	81	757	6952	142
STARCRAFT GRANDMASTER	MEAN	28	177	61	215	727	7566	133
	MAX	28	179	61	363	848	7566	133
ATARI-NET	BEST MEAN	25	96	49	101	81	3356	< 1
	MAX	33	131	59	351	352	3505	20
FULLY CONV	BEST MEAN	26	103	45	100	62	3978	3
	MAX	45	134	56	355	251	4130	42
FULLY CONV LSTM	BEST MEAN	26	104	44	98	96	3351	6
	MAX	35	137	57	373	444	3995	62

[7]

Learning from Replays - Value Predictions

- Supervised Learning



[7]

Learning from Replays - Policy Predictions

	TOP 1 ACCURACY			TOP 5 ACCURACY		
	ACTION	SCREEN	MINIMAP	ACTION	SCREEN	MINIMAP
ATARI-NET	37.8%	1.2%	19.8%	87.2%	2.9%	55.6%
FULLYCONV	37.9%	9.5%	25.7%	88.2%	18.5%	62.3%
ARFULLYCONV	37.7%	10.5%	25.9%	87.4%	22.1%	62.7%
RANDOM	4.3%	0.0%	0.0%	29.5%	1.0%	1.0%

[7]

QUICK REVIEW

**SC2LE
Overview**

Mini Tasks

**Supervised
Learning**



VOID



THE MASTERY


Part II

What has happened? – A new star is born

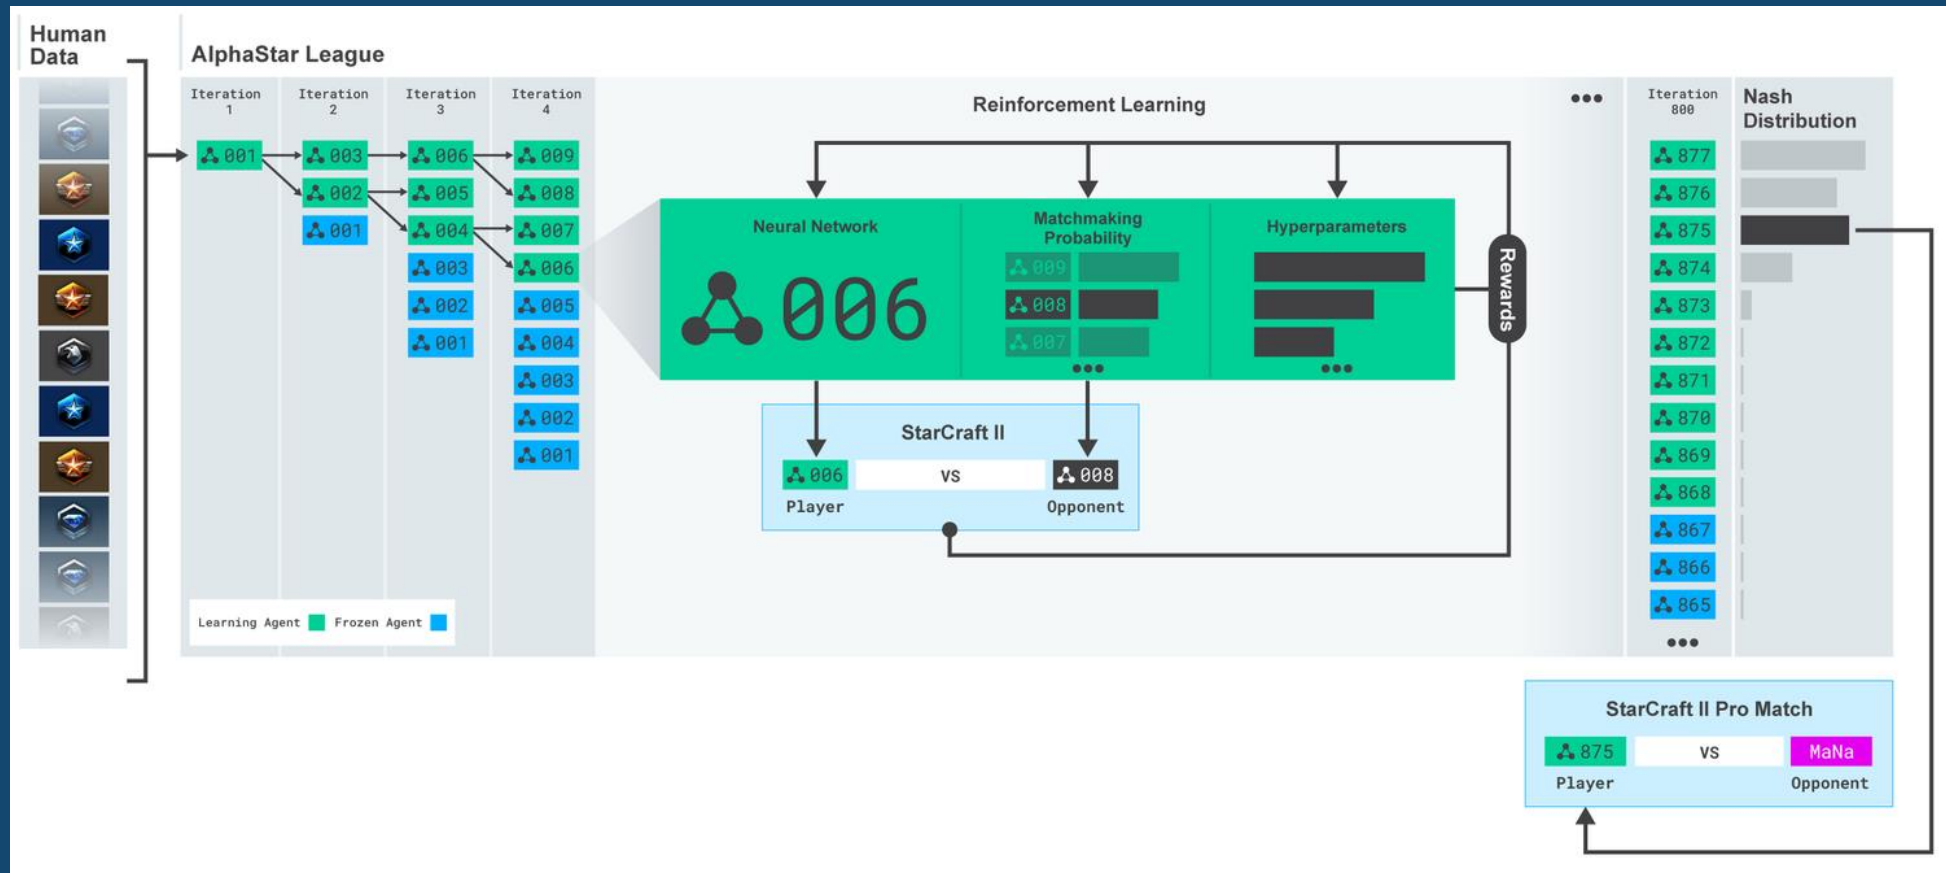
- December 10th 2018: AlphaStar beats the best DeepMind Starcraft player
- December 12th 2018: AlphaStar beats Dario “TLO” Wunsch, a Pro Starcraft Player
 - BUT: TLO plays Zerg normally
- December 19th 2018: AlphaStar beats Grzegorz “MaNa” Komincz, a Pro Starcraft Protoss Player

AlphaStar – What is inside? [10]

- Deep LSTM Core: sequence modelling, natural language processing (NLP) [14]
- Transformer Architecture: Attention mechanism, parallel computation [15]
 - Pointer Network: Use attention as pointer to input [16]
- Auto-regressive Policy: Use previous observations for next prediction [7]
- Centralised Value Baseline instead of a Multi-Agent system [17]




AlphaStar Training




[10]

AlphaStar League – MMR



[10]


Evolving Strategies



[11]




[12]



[13]

Nash distribution in AlphaStar League




[10]

Training the League

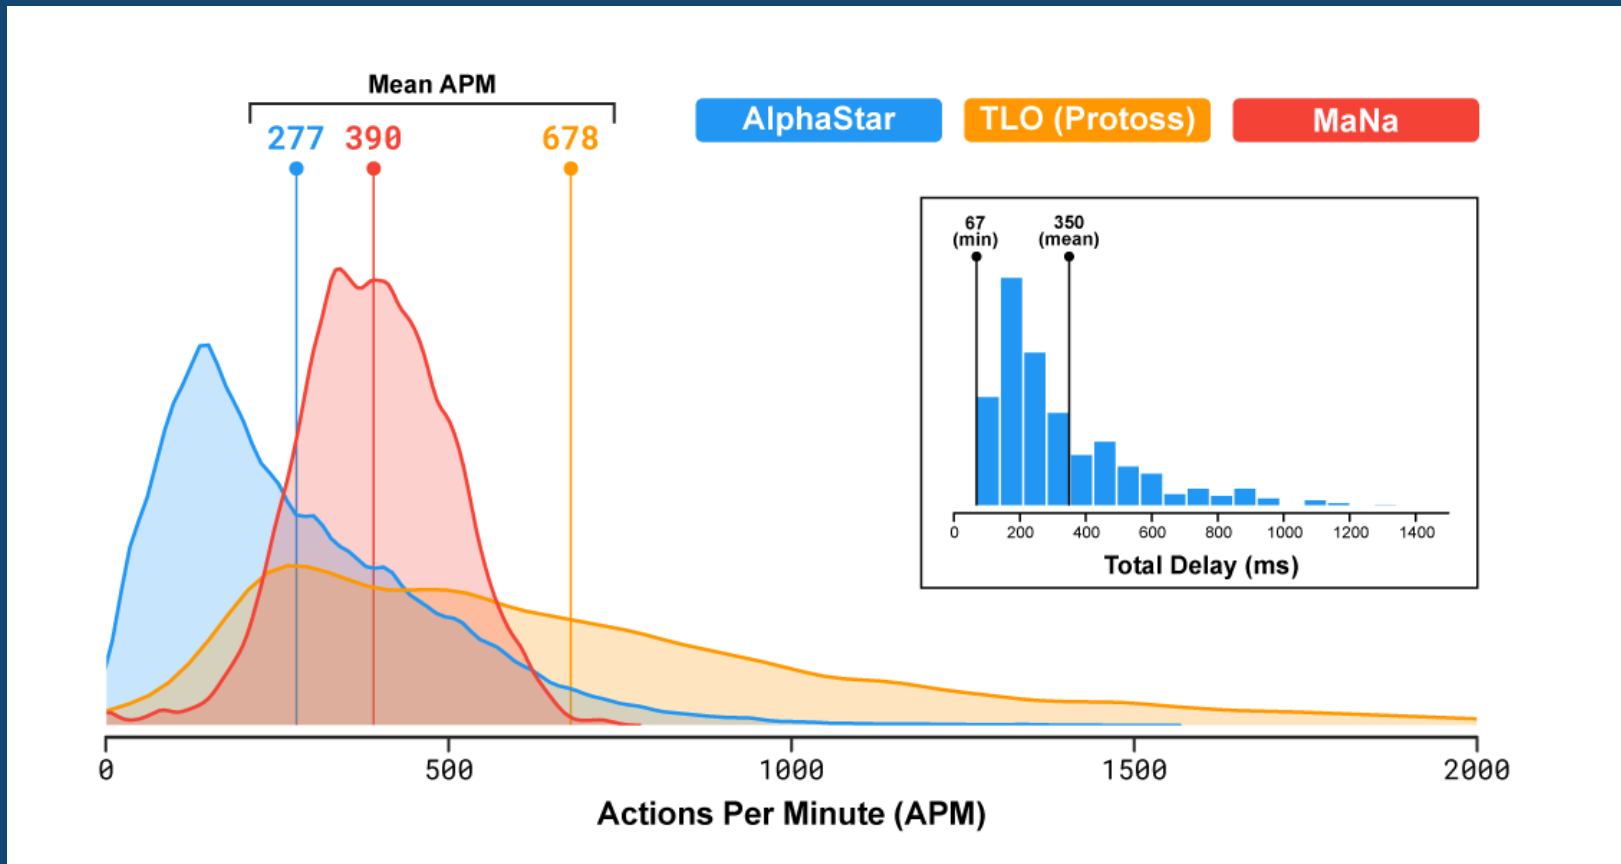
- 14 days of training
 - 16 TPUs per agent
- => up to 200 years of Starcraft play per agent

Example




[10]

Comparison to Human Play



[10]

Comparison to Human Play



[10]

NEWS!

11.07.2019

- Announced yesterday: AlphaStar will play online in competitive ladders in Europe [18]
 - All races (Terran, Zerg, Protoss)
 - Camera-like view
 - Anonymously 😞
- => Go play Starcraft (It's free!)
- Future: AlphaStarZero?

More about AlphaStar

AlphaStar – Inside Story [19]



AlphaStar Demonstration [20]





THANK YOU FOR YOUR
ATTENTION!

ANY QUESTIONS?

References

- [7] <https://arxiv.org/pdf/1708.04782.pdf>
- [10] <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>
- [14] <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.676.4320&rep=rep1&type=pdf>
- [15] <https://arxiv.org/pdf/1706.03762.pdf>
- [16] <https://papers.nips.cc/paper/5866-pointer-networks.pdf>
- [17] <https://www.cs.ox.ac.uk/people/shimon.whiteson/pubs/foersteraaai18.pdf>
- [18] <https://starcraft2.com/en-us/news/22933138>
- [19] <https://www.youtube.com/watch?v=UuhECwm31dM>
- [20] <https://www.youtube.com/watch?v=cUTMhmVh1qs>

Image Sources

- [1] <https://logonoid.com/starcraft-2-logo/>
- [2] <https://www.youtube.com/watch?v=CXe06EsUexQ>
- [3] <https://www.kotaku.com.au/2015/09/even-the-koreans-think-starcraft-2-is-too-hard/>
- [4] <https://www.youtube.com/watch?v=UuhECwm31dM>
- [5] <https://siliconangle.com/2016/11/04/google-deepmind-to-use-the-messy-world-of-starcraft-for-ai-research/>
- [6] <https://www.businessinsider.de/david-silver-the-unsung-hero-at-google-deepmind-2016-3?r=US&IR=T>
- [7] <https://arxiv.org/pdf/1708.04782.pdf>
- [8] https://starcraft2.4fansites.de/galerie_6_1009.html
- [9] https://www.freepik.com/free-icon/stopwatch_739036.htm
- [10] <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>
- [11] <https://starcraft.fandom.com/wiki/Stalker>
- [12] <https://www.deviantart.com/ghostnova91/art/Adept-Placeholder-551517749>
- [13] <https://www.youtube.com/watch?v=Ej0aXs2xJIA>