



Simulating Teamfight Tactics

Using Deep Learning for Fast Reinforcement Learning AI Training



Ran Cao

Data Scientist, Riot Games





Agenda

Intro to Teamfight Tactics(TFT)!

A round based strategy game!

Simulating TFT

Powered by deep learning!

Training AI

How we speed up the training

Value & Limits



Story

Teamfight Tactics is fun!

We want to develop AI to help with the design process

Challenges and decisions





TEAMFIGHT TACTICS™




TEAMFIGHT TACTICS™

Round Based
Strategy Game
League Of Legends Universe





Draft Your Unique Team Of Champions



Your Team



You



Draft Your Unique Team Of Champions







































































Your Team





TFT Set 1 Champion Roster

Tier 1 (1 gold)	 Tier 2 (2 gold)	 Tier 3 (3 gold)	 Tier 4 (4 gold)	 Tier 5 (5 gold)
 	 	 	 	
 	 	 	 	
 	 	 	 	
 	 	 	 	
 	 	 	 	
 	 	 	 	
 	 	 	 	





5-6





5-6

- 6 Heart 6
- 1 Corrupted 1
- 3 Star Guardian 3
- 2 Frankster 2
- 2 Argle 2
- 2 Mascot 2
- 2 Ox Force 2
- 2 Spelllinger 2
- 1 Threat 1
- Gadgetron 173

- 80
- 69
- 14
- 18
- 10
- 0
- 0
- 0

RIOT GAMES

Lvl. 9

4

2

6

Buy XP

4

Refresh

2

LaserCorps

Sureshot

Senna # 3

Anima Squad

Ace

Miss Fortune # 4

Threat

Bel'Veth # 4

MechaPRIME

Brawler

Jax # 3


Star Guardian

Mascot

Heart

Yuumi # 2

11



Each Round You Fight Another Player's Team





Why AI is a Good Fit For Teamfight Tactics

Playtesting is hard.

- Situations in one game is not easily replicable
- Needs 8 evenly skilled players for each game
- Every few months, **TFT** will update to a fresh new set





Goals

Train a TFT AI that can play the game at high skill level

- Method should be practical
- Training should be fast enough for the content cycle





Challenge #1

TFT game server integration

- **TFT** game server doesn't have the foundation for an AI system.
- A full game of TFT takes 40 mins, we want one AI game to end in seconds.





(TF)T-MINUS EIGHTEEN WEEKS: THE STORY OF TFT (PART 1)

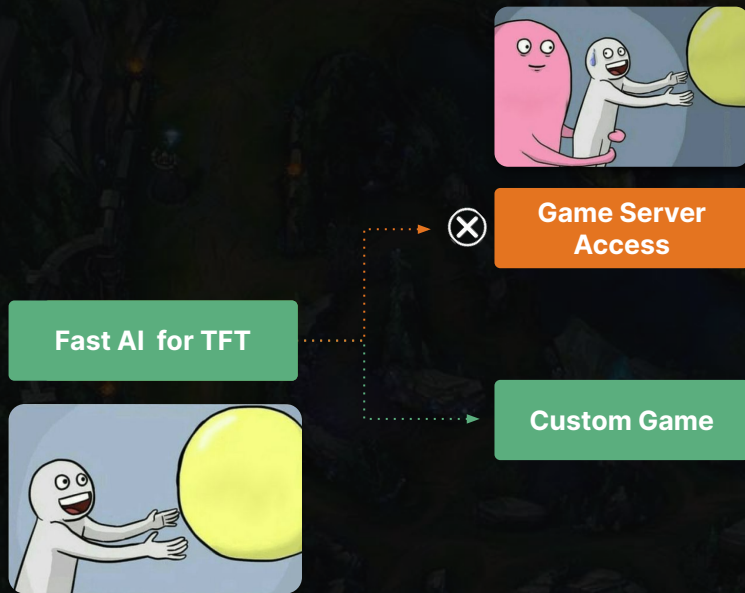
What do you get when you add 12 devs, eight weeks, and League of Legends together? TFT, obviously.



Riot Stimhack

DATA SCIENTIST — RAN CAO

Ran develops machine learning features for Teamfight Tactics. She knows your MMR, she gives you LP, and sometimes she matches you against the strongest player in the lobby.



Custom Python Game

Current Round: 4-1 Current Phase: Phase.DEPLOY Current Player: 0 Swap Mode

Name: player0
Level: 7
XP: 6/56
Gold: 3
Health: 50.0

OutId-1
OutId-2
Scalescorn-2
Bruiser-2

Items:

Bench:

Summons:

Carousel:

Buy Menu:

Level Up: 0.00
Reset: 0.00
End Turn: 0.00

player0: 50.0
player1: 35.0
player2: 74.0
player3: 63.0
player4: 92.0
player5: 43.0
player6: 57.0
player7: 46.0



We need a solution to resolve teamfight and predict outcomes





Input states

Blademaster - 6
Celestial - 2
Reaver - 2




Damage: -12

Vanguard - 4
Infiltrator - 2



Output damage distribution





Generalized Image

Blademaster - 6
Celestial - 2
Reaver - 2



Damage: -12

Vanguard - 4
Infiltrator - 2



Convolutional Network

A hex \rightarrow A pixel

A champion embedding \rightarrow R,G,B

CNN Input:

						17
					2	
		67	98	88	143	34
		71	122	78	82	
90		75		95		



Generalized Sentence

Blademaster - 6
Celestial - 2
Reaver - 2



Damage: -12

Vanguard - 4
Infiltrator - 2



Transformer Encoder


An occupied hex → A word
Hex position → word position

Champion "words" encoding:
[90,75,95,71,122,78,82,67,98,77,143,34,2,17]

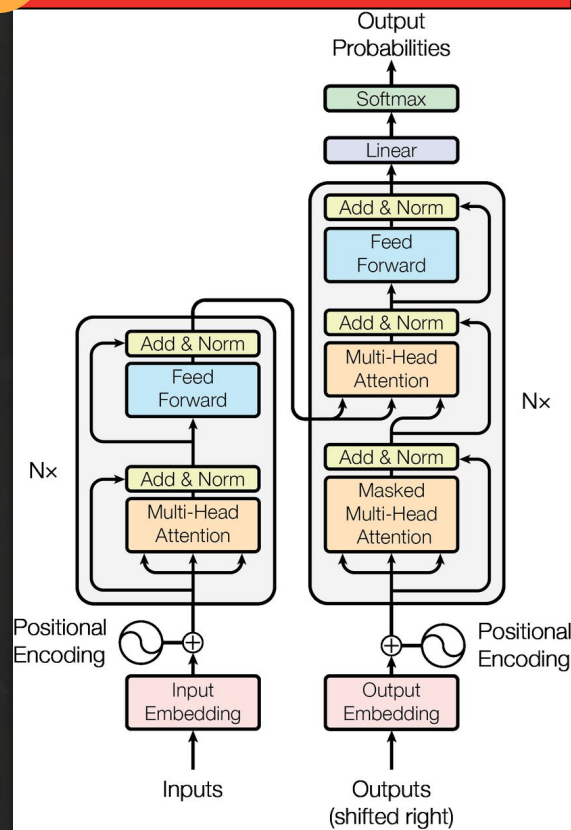
Champion position encoding:
[1,3,5,17,18,19,20,54,53,52,51,50,37,29]



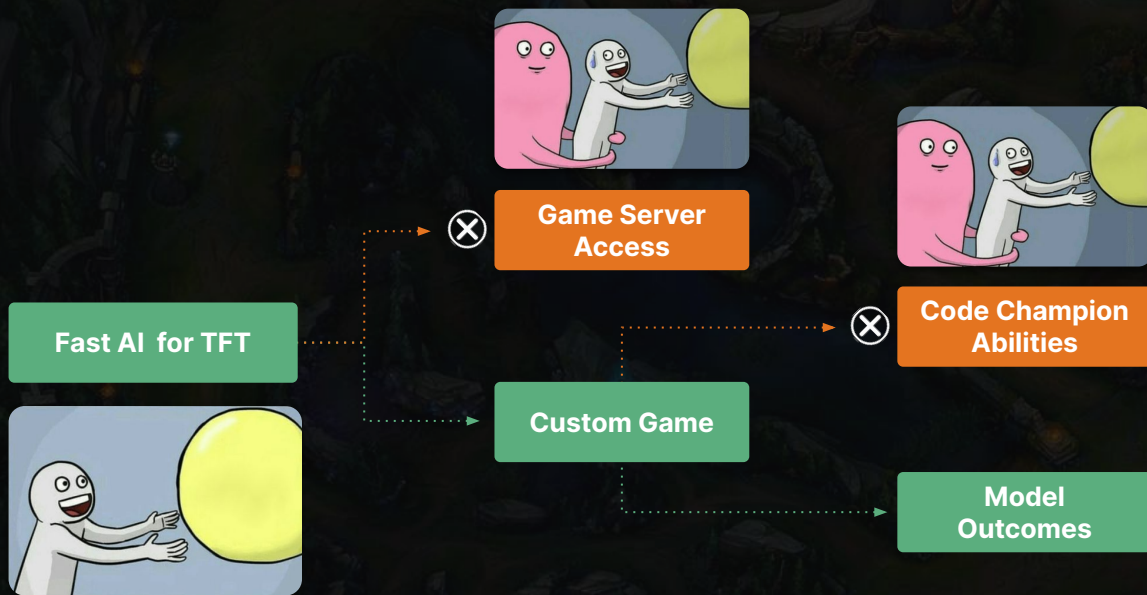
Models (CNN+Transformer Encoder)



Attention is all you need



CNN < Transformer < CNN + Transformer






How to train a self-learning AI ?



Reinforcement Learning Effort within Gaming



OpenAI Five Defeats Dota 2 World Champions

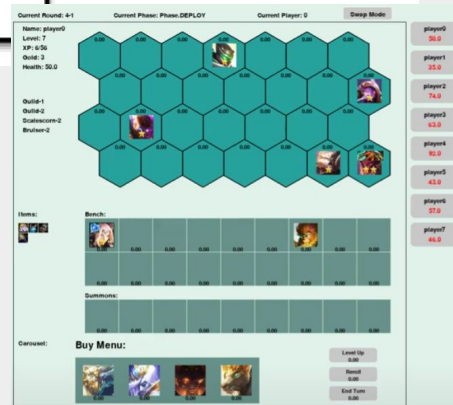

OpenAI Five is the first AI to beat the world champions in an esports game, having won two back-to-back games versus the world champion Dota 2 team, OG, at Finals this weekend. Both OpenAI Five and DeepMind's AlphaStar had previously beaten good pros privately but lost their live pro matches, making this also the first time an AI has beaten esports pros on livestream.



How does it work?

Reinforcement Learning

Reward





Start Simple





**Not just start simple,
set milestones.**

It is going to be a long journey.



What marks a new milestone?

A significant increase over environment complexity:

- More states
- More actions

Examples:

- A new skill expression
- More players
- Increased game length





TFT AI Training Milestones

M1. 10HP 1v1

M2. 50HP 1v1

M3. Unit Deployment

M4. Unit Positioning

M5. 100HP 1v1

M6. 8-player Free-for-all

M7. Draft & Equip Items





TFT AI Training Milestones

M1. 10HP 1v1

M2. 50HP 1v1

M3. Unit Deployment

M4. Unit Positioning

M5. 100HP 1v1

M6. 8-player Free-for-all

M7. Draft & Equip Items





Challenge #2

Deployment Phase






Highlighted Challenges:

Longer feedback loop


More complex credit assignment


A Round




**Combat
HP changes**

Next Round




**Game End
Get Reward**
**Combat
HP changes**

Last Action

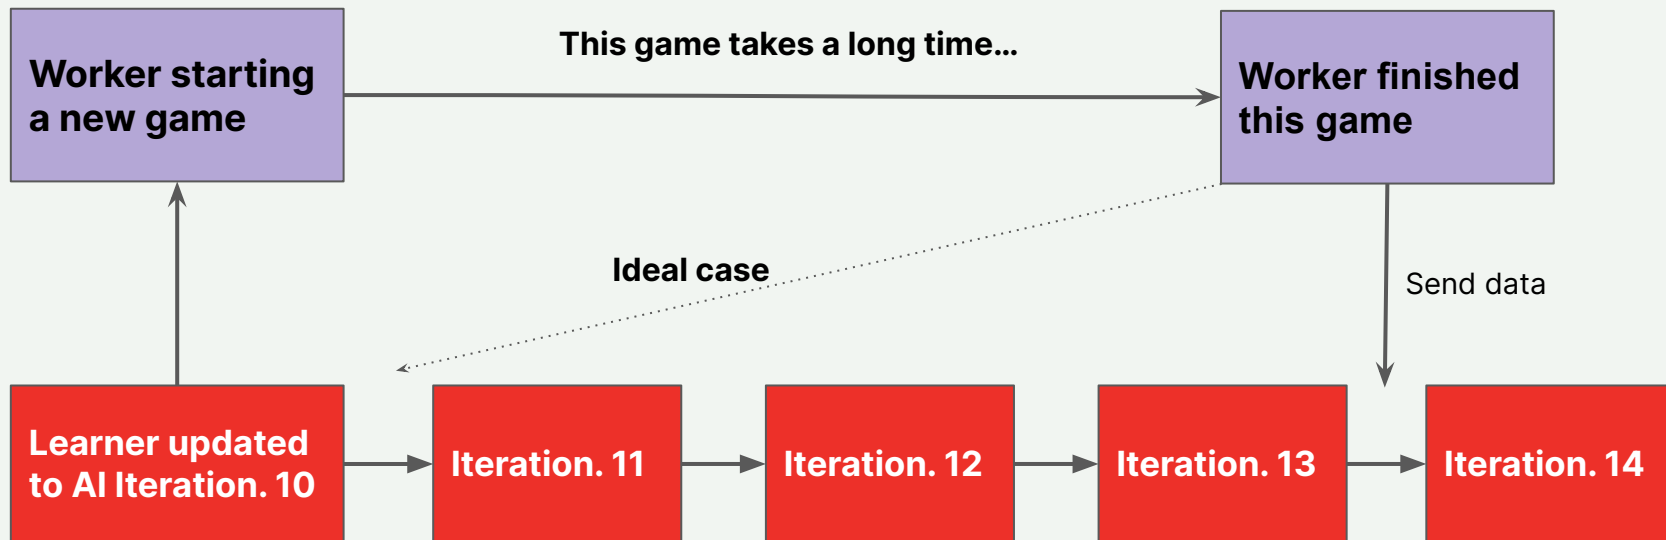


Highlighted Challenges:

Longer feedback loop

More complex credit assignment

Increased Off-policy degree





Instead of having **1 AI** playing the entire game



Only goal: Win the game!

We decide to let **2 AIs** playing different Phases as one player



Agent.1 goal: Win the game!

Agent.2 goal: Win this round's combat!





Shop AI Action Space

Gold Actions

- End Turn
- Reroll
- LevelUp
- Buy unit
- Sell unit

Item Actions

- Equip Item
- Combine Item

Reward



Deployment AI Action Space

Deploy and Position Actions

- Deploy (my 3rd unit)
- Choose a board position for the selected unit

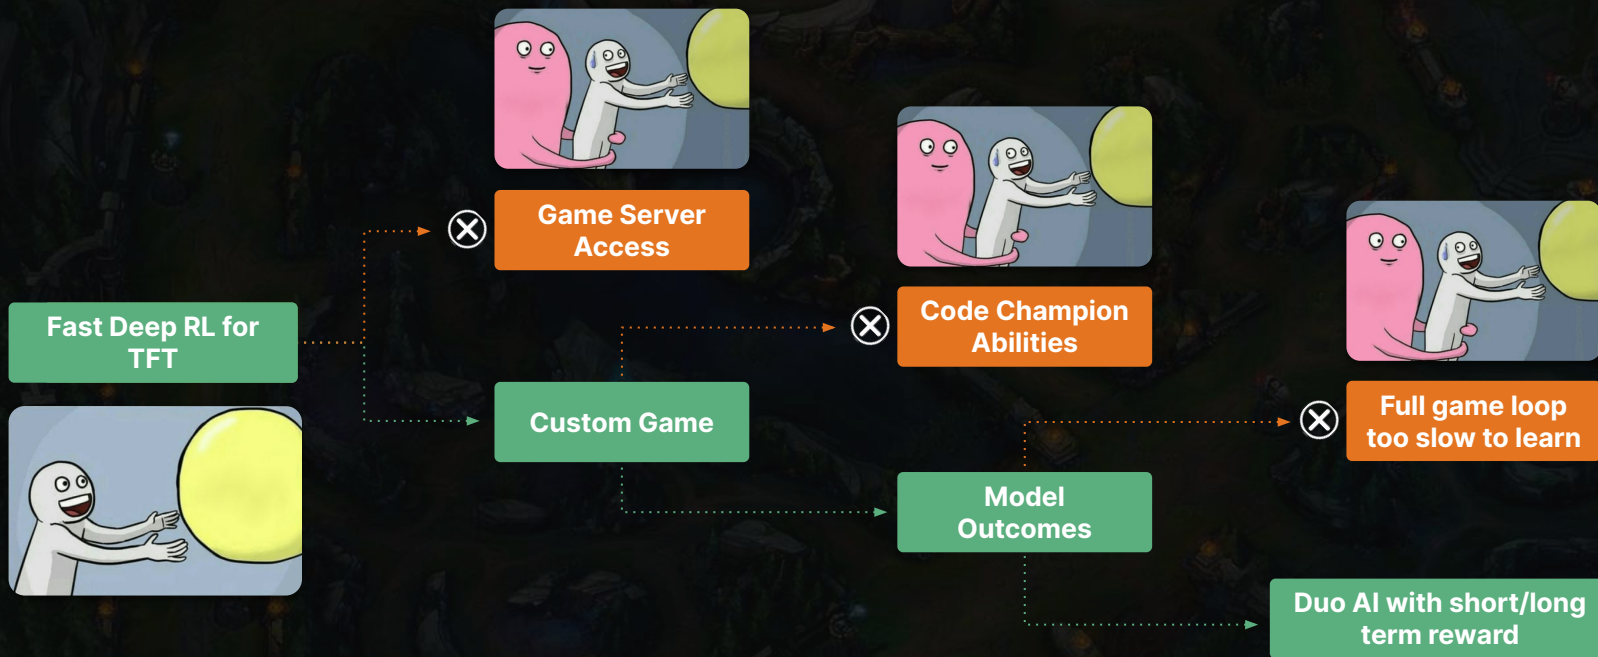
Reward:
Round end
combat damage
outcome





Shop Agent Learns Much Faster Than Full Agent





**Now I have some
agents, what's next?**



Balance testing

TFT 12.17B-Patch notes

TRAITS

- Cannoneer 5th attack cannon shot damage: 175/350/550% \Rightarrow 150/300/475%
- Darkflight bonus Health 600/800/1000/1800 \Rightarrow 500/750/1000/1800


UNITS

- Aphelios Attack Damage: 65 \Rightarrow 60
- Aphelios Attack Speed: 0.8 \Rightarrow 0.75
- Jayce Mercurial Justice Knock Up duration: 1.5 \Rightarrow 1.25
- Olaf starting Mana nerf: 50/100 \Rightarrow 30/100
- Olaf Attack Damage gained upon death: 5 \Rightarrow 4
- Pantheon Bugfix: Pantheon is now Mana locked for the duration of his spell
- Pantheon Aegis Assault Attack Damage ratio 280/300/400% \Rightarrow 265/280/380%
- Shyvana Dragon's Descent Stun duration: 1.5 \Rightarrow 1.25
- Xayah Attack Damage: 75 \Rightarrow 70


AUGMENTS

- Knife's Edge Attack Damage: 20/35/50 \Rightarrow 15/25/40

Certain units got detected within hours



Units that's harder to play takes days to detect





Reward Shaping Experiments that you can't do on real players

Rank progression and Strategy Depth


Sweet Spot

Winner takes all

1st - 8th reward in a ladder

Top 4 wins

...







Limitations

Keep up with the live game is hard

Communicating your learning progress with stakeholders can be difficult.

Still need to integrate and fine tune in the real game server.





THANKS!



Ran Cao

Data Scientist, Riot Games

RCAO @ RIOTGAMES.COM