

Artificial Intelligence: A report on the multi-agent platform Arena

Carolina Proietti

Julian Marks

Juan Diego Arango

1. Background

The paper is about a general evaluation platform for multi-agent intelligence that may be implemented with different games. Understanding the importance of this area, the authors show that there is not a general evaluation platform for multi-agent intelligence, all platforms that exist aren't compatible with multi-agent settings or are limited to specific games. That's why the authors present Arena, a general platform for evaluating multi agent intelligence, with 35 games of diverse logic and representation.

Arena is a building toolkit for multi-agent intelligence, that includes five different multi-agent reward schemes and a configurable social tree. Authors explain how the emergence of innovation is becoming a really hot topic in AI. Agents in Artificial Intelligence already can solve tasks given by programmers and users in an impressive manner and may also exceed human capabilities in many fields, but often they fail to innovate and can not come up with creative solutions to given tasks. The authors hope to tackle this problem with their special take on multi-agent programming: "Each agent is creating an environment for the others, and improving itself implies creating new problems for the other agents" [Arena Publications-Latest pdf]. This is really important because in multi agent learning, strategies for the agents aren't defined by the programmer, each agent creates its own strategies based on experience and rewards.

This kind of learning is a very good approximation to the real world and that's the reason why its thriving.

In the second chapter the Paper is about the graphical engine used in Arena, which is the world-leading gaming engine Unity. Comparisons to other engines are made, followed by an in-depth chapter about the implementation of the reward and social tree configurations as well as the learning process of the agents. Following this the authors show how to conduct experiments using Arena, talking about stochasticity and evaluation techniques. After that comparisons to other, similar papers are made followed by a conclusion about the whole platform.

The developers of Arena expect to help researchers create new games and be able to test them in a standardized way with various configurations.

2. Contributions

As said previously, innovation is becoming a very important topic for artificial intelligence and one of the most promising paths towards such a vision is learning via social interaction. Multi-agent systems can solve problems that are difficult or impossible for an individual agent or a monolithic system and, despite enormous and sophisticated strategies are invented, according to Arena team there isn't yet a functional method to learn through the collective interplay of agents. According to Arena team "In multi-agent learning, how the agents should beat the opponents or collaborate with each other is not defined by the builder of the environment, e.g., the inventor of the ancient 'Go' never defines what strategies are good." [Arena Publications-Latest pdf].

Specifically, Arena provides: (1) 27 new games that are yet not studied in the community; (2) 8 games, of which the basic logic is inspired by other lectures but equipped with realistic rendering effects, physics engine and other features, such as extensibility to other social paradigms; and (3) interface to the popular stand-alone

domain StarCraft. Many of the in-built games are either new or contain yet unsolved problems.

The authors also supplies a python implementation of five deep multi-agent reinforcement learning baselines among different kinds of agents and other tools. Thanks to that it's possible to train in different training schemes in order to evaluate them in different scenarios. They involve both competitive settings, like fights against each other until just one agent remains and collaborative ones like being in a team and having to organize against the other team until just one team is left. This is aimed to allow to understand in depth how the several agents improve and react to the context "in a way never done before". The diversification of the set games provided allows the users and researchers to explore the many problems of multi-agent intelligence, giving them the possibility to approach multi-agent AI "in a stable and uniform way".

In order to proceed to the more interesting high-level-intelligence tasks there are several pre-implemented algorithms for basic intelligence, for example movement skills of the agent and the will to stay alive in a competitive setting against other agents. It seemed to be a long-time burden for developers to have to build the basic intelligence every time before being able to take care of more advanced tasks. Furthermore, the reward schemes of the agents can be easily changed in such a way that different social relationships can be implemented, thus simplifying the creation of complex social trees.

Finally, there is a necessity for competitive agents to evaluate against human players, and also a research trend for collaborative agents to team up with human players. Thus, Arena provides a gaming interface for humans, so that a human player can take the place of any agent in the game scene.

The engine behind Arena is Unity, the authors compare it to other widely known engines and point out the advantages that Arena offers to developers, like enabling multi-agent

intelligence with simple customization and up-to-date graphics with advanced rendering, lighting and realistic physics and a vivid community. They also emphasize the wide range of game genres offered, from real-time strategy to shooters and also detach from some dated softwares like DeepMind, PsychLab and Malmo.

3. Software Review

We spent a bunch of time trying out Arena, learning how it works and understanding the different scenarios it has to offer. This is our software review:

Requirements:

Arena environment doesn't have a lot of requirements, it only requires Unity, therefore Arena requirements are Unity ones:

OS: Windows 7 SP1+, 8, 10, 64-bit versions only; Mac OS X 10.9+.

CPU: SSE2 instruction set support. GPU: Graphics card with DX9 (shader model 3.0) or DX11 with feature level 9.3 capabilities.

Installation:

Installing Arena was really simple, we installed Unity first (just downloading it from the web page) and then we cloned the two Arena repositories. We opened one of the folders in Unity and Arena was launched.

Usability:

Arena is built on unity so one must know how to use unity to use Arena, it is like a native extension of unity. We didn't know Unity so we had to learn the basics to try and test Arena. Unity is not easy, so we spent a long time there. Having an idea of Unity, we tried out the different arena prefabs such as agents, backgrounds and even games. Our idea was to try to make a thinking group of agents play some Arena prefabricated game

and make some analysis on their performance and stuff but it was impossible. Launching an already built game in Arena is not complicated, and we were able to move the characters and stuff by ourselves. Understanding completely how the social tree works, especially try all of the features and the setting of the Arena node wasn't too easy, but thanks to a tutorial we were able to learn the most important commands and we have tried both the collaborative and competitive mode. The problem was making the Agents play on their own. To be able to do that, some pretty complicated python modeling was needed. Arena offers this "brains" that are deep learning networks for the agents to learn and think. The difficulty is that those brains have to be trained with python models. We tried our best but there was a lot of requirements that needed to be known to be able to teach the brains so we couldn't do it. Only the part of the documentation to try and do that kind of stuff was there, but it was really complex and the tutorials about how to proceed after importing the models were non-existent. Even though we weren't able to make agents play against each other in a game, we still wanted to try other Arena features. We created a simple two player game (two players in a square board) and tried different reward schemes of Arena and it was interesting. We could put different rewards such as distance from a target, or longest time in the board or shortest time in the board. We could manually move players and see how the rewards were changing and it worked really good.

4. Conclusions

According to our experience, we think that Arena is a very innovative platform in its field. Thanks to Unity we had an easy experience once we had learned the basic commands, so the approach to the program is really intuitive. Furthermore, with more in-depth use of the software and a greater knowledge of the part concerning the programming of agents, Arena constitutes an important step for the research and it could really achieve the goal of innovation to which it aspires. We think Arena is an amazing software and

offers really good opportunities for research, but it's designed for people who know a lot about Unity, Python and the usage of those languages for modeling multi-agent environments and reinforcement learning, for students it is not very useful without a very broad background. Still, all in all, this contribution is an amazing accomplishment regarding that Arena was created collaboratively by students of universities both in China and the United Kingdom.