

Starcraft II "Lite" Environment For Multi-Agent Reinforcement Learning Research

Adam Michalski



Master of Science
School of Informatics
University of Edinburgh
2022

Abstract

With artificially intelligent agents entering more and more areas of our daily lives, we need methods to ensure we are only using the best algorithms to guarantee our safety and high quality of life, and there is a clear lack of standard benchmarks for Multi-Agent Reinforcement Learning (MARL) algorithms. The Starcraft Multi-Agent Challenge (SMAC) claims to be a good candidate for a standard MARL benchmark, but it is fully based on top of a heavy, closed-source computer game Starcraft II. This makes it completely closed to extension for researchers not familiar with the game, and requires the knowledge and use of other proprietary tools specific to the game for any meaningful contribution to the challenge. We introduce SMAClite – a challenge based on SMAC that is both completely decoupled from Starcraft II and open-source, along with a framework which makes it possible to create new content for SMAClite without any special knowledge. We conduct experiments to show it is equivalent to SMAC in many ways, and it outperforms SMAC in both speed and memory.

Research Ethics Approval

This project was planned in accordance with the Informatics Research Ethics policy. It did not involve any aspects that required approval from the Informatics Research Ethics committee.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Adam Michalski)

Acknowledgements

I am grateful to everyone who supported me on this adventure, whether it be by offering words of support, letting my mind rest a bit by accompanying me on unrelated side-quests when I needed them, or covering for me in various areas when my mind was preoccupied with matters it considered more important.

Special thanks are owed to my dissertation co-supervisor Philippos, who was always present and patient to answer all sorts of questions that I came up with, and to my managers Ewa and Robert at Amazon who were unbelievably accommodating to any university-related requests I might have had throughout this whole journey, even though I do not believe there was any way these requests were beneficial for business.

I would like to thank my significant other Christina for an incredible amount of support and for insisting to me that I am capable of succeeding, even when I was sure I was not.

Table of Contents

1	Introduction	1
1.1	Motivation	1
1.2	Objectives	2
1.3	Related Work	3
1.4	Report Structure	4
2	Background	5
2.1	Multi-Agent Reinforcement Learning	5
2.2	MARL algorithms	6
2.2.1	Independent Learning	6
2.2.2	Centralized Training, Decentralized Execution	7
2.3	SMAC	7
2.3.1	Overview	7
2.3.2	Actions	8
2.3.3	State	9
2.3.4	Observations	10
2.3.5	Rewards	10
2.4	Optimal Reciprocal Collision Avoidance	11
2.4.1	Overview	11
2.4.2	ORCA Methodology	11
3	Project Methodology	13
3.1	Environment implementation method	13
3.2	Base framework	14
3.2.1	Scenario definition	15
3.2.2	Unit definition	16
3.3	Unit command types	17

3.4	Game loop	18
3.4.1	Target clean-up	18
3.4.2	Velocity preparation	19
3.4.3	Velocity adjustment	20
3.4.4	Game step execution	20
3.5	Custom ORCA implementation	21
3.6	Opponent AI behaviour	22
4	Experiments	24
4.1	Agent learning curves	24
4.2	Behaviours observed in agents	28
4.3	Cross-environment 0-shot performance	29
4.4	Environment Performance Benchmark	30
4.4.1	Environment speed	30
4.4.2	Environment memory usage	31
5	Conclusions	32
5.1	Future Work	32
5.2	Summary	33
	Bibliography	34
A	SMAClite framework examples	40
A.1	Example of a valid scenario file	40
A.2	Exmaple of a valid unit file	41
A.3	Standard unit types	41
A.4	Terrain presets	42
A.5	Standard scenarios	42
A.6	Unit attributes	42

Chapter 1

Introduction

1.1 Motivation

As we continue to make advancements into artificial intelligence research, it inevitably makes its way into our daily lives. Examples of autonomous agents popular in recent times range from robotic vacuum cleaners (e.g. [1]) and self-driving cars (e.g. [2]) to robots making lives easier from behind the scenes, such as warehouse optimization robots (e.g. [3]). With all this attention in research, a natural need arises for standardized benchmarks for the various types of artificial intelligence (AI) models, to always make sure we are using the best technology available to make our everyday lives easier.

Multi-agent reinforcement learning (MARL) is a branch of machine learning dealing with multiple autonomous AI entities – usually called *agents* – existing in the same space. In this work, are interested in *cooperative* agents, ones that work together to accomplish some goal – like those one would use in self-driving cars to avoid collisions and reach their destination quickly. Interestingly, there is currently no consensus in the research community about what a standard benchmark for this type of agents should be. Just by looking at a recent benchmarking paper [4] we can count 5 different benchmarking tasks.

The Starcraft Multi-Agent Challenge (SMAC) [5] claims to be one candidate for such a standard benchmark. It is built on top a real-time strategy computer game Starcraft II (SC2), and makes use of an API – an interface between the game and the AI agents – made available by Vinyals et al. [6]. It presents a mini-game where each agent controls a single combat unit (i.e. a single soldier) in one of several available battle scenarios against an enemy team controlled by the game’s built-in AI. Its authors claim that SMAC is a good benchmark because it presents a challenge where the solution

is not straightforward – in most of the scenarios the most obvious strategy of running forward and attacking is not good enough, and will result in a quick loss due to the enemy army having better units or more numbers.

While the idea is promising, we can spot several problems with SMAC if it is to become a universally accepted benchmark. The biggest issue we see is that it remains highly inaccessible for people unfamiliar with the game it is based in. This manifests in many ways, e.g. using terms from the game such as "zergling" or "marine" without further elaboration. It simply might not be obvious to those of AI researchers who do not have experience with SC2 what the difference between the two is. To create custom scenarios or units for SMAC, one is required to use the official Starcraft II map editor, requiring people to learn an unusual and proprietary tool, and put in a lot of effort for a single benchmark.

On top of that, SMAC itself uses SC2 as its key dependency, requiring a large-sized (ca. 3.7 GB) download and a complicated setup process for any training or inference, not to mention running SC2 alongside SMAC at all times, consuming extra CPU and memory resources. This is made even worse by the fact that SMAC uses only a subset of the SC2 features – a lot of the required downloadable and computational resources are simply redundant, and due to that training agents on SMAC is more expensive than necessary for the task it presents.

1.2 Objectives

Our main goal for this project is to **implement a challenge similar to SMAC, but completely decoupled from the Starcraft II dependency, and prove it preserves the challenging aspects of SMAC**. We name this challenge Starcraft Multi-Agent Challenge lite (**SMAClite**). We also wish to make the battle scenarios and units as easy to modify as possible (also allowing easy creation of completely new ones), all while keeping the installation process as simple as possible. We wish to preserve the outer interface of SMAC, only changing the inner workings to match our goal, to allow AI developers to reuse their code for handling SMAC with minimal modifications.

Roughly, our work can be divided into three milestone objectives. First of all, we aim to implement the challenge itself, together with the opponent AI and a simple 2D rendering solution allowing to visualize what is happening inside. We wish to maintain parity with Starcraft II as far as it is necessary to maintain parity with SMAC, and omit anything else. One of the goals we keep in mind while implementing SMAClite

is environment performance, doing our best to keep it as fast and as cheap to run as possible.

Second of all, we wish to make the training of multi-agent reinforcement learning algorithms with SMAClite as frictionless as possible. For this reason, we aim to integrate it with the ePyMARL framework for training and running MARL models, introduced by Papoudakis et al. [4]. With this accomplished, training a model in SMAClite should be as simple as running a single terminal command.

Finally, we will run experiments using the trained models, to verify the deliverables of our two previous milestones are sound. Our experiments will include quantitative analysis by comparing the reward (standard MARL metric, see section 2.1) achieved by the various MARL algorithms in SMAClite, as well as qualitative analysis, looking at the combat strategies employed by the agents on a case-by-case basis, and verifying they do indeed outsmart the handwritten enemy AI. On top of that, we wish to take the models trained on SMAClite, and put them inside the original SMAC environment without any further training, to see how much potential for transfer learning there is between the environments.

1.3 Related Work

It is important to note that popular benchmarks, accepted by the research community as the standard, do exist for the single-agent variant of reinforcement learning. MuJoCo [7], based on the framework by Todorov et al. [8], is a set of tasks for robotic manipulation in a 3D world. Arcade Learning environment [9] is a suite of 55 classic arcade games with an interface that allows AI agents to attempt to master them. Both of them are widely used as reinforcement learning benchmarks.

There are several non-Starcraft II multi-agent benchmarks worth mentioning. Multi-Agent Particle Environment [10] places the agents in a 2D environment and has them accomplish various physical tasks such as moving into specific positions or pushing objects around, also allowing for some limited communication between agents. The Hanabi challenge [11] puts the agents in the cooperative board game Hanabi, which only allows for limited, strictly formatted communication. The authors of the MARL benchmark paper [4] also make available two more benchmarks, both based in simple 2D worlds: Level-Based Foraging, where agents have to collect food placed in the world, and Multi-Robot Warehouse, where agents have to locate and deliver specific items to specific shelves in a warehouse.

One other multi-agent challenge based on a modern computer game is the OpenAI Five project [12], which put agents inside a full five versus five match of a strategy game Dota 2, showing impressive scale of the game and the trained models.

Training autonomous agents in Starcraft II started becoming popular upon the publication of its API [6], with the most popular result being the AlphaStar model, said to be capable of beating 99.8% of officially ranked human players at the full game of Starcraft II [13]. Our work is mostly based on the work of the creators of SMAC [5], of which our project can be considered major rewrite with added accessibility features and improved performance. Since its publication, SMAC has been used as a benchmark for autonomus agents in numerous works (e.g. [4, 14, 15]), with its popularity being a sign that it does fill a niche – it is our hope that our project will help it become more widely adapted.

There are several projects that are tangentially related to ours, since they also aim to improve SMAC, but take a different approach. SMACv2 [16] seeks to make the challenges more difficult by randomizing the team compositions and starting locations, and SMAC+ [17] increases its difficulty by adding scenarios requiring the agents to solve multi-step challenges including mazes, making use of terrain and line of sight restrictions. Both of these present interesting additions to the SMAC paradigm, however neither addresses the issues we wish to tackle. We believe the additions introduced by them are good ideas for future developoment of SMAClite, but ideally, in the scope of this project, we would like to maintain the challenge’s difficulty on the same level as the original.

1.4 Report Structure

We structure our report as follows: in chapter 1 we introduce the problem and briefly describe what we wish to accomplish. Then, in chapter 2 we describe in detail the preliminary knowledge necessary to understand our work, and in chapter 3 we provide a thorough specification of our project. We continue in chapter 4 with a description of the experiments used by us to verify our implementation is sound, as well as what conclusions we can draw from them. We conclude in 5 with ideas for future work related to our project and a summary of our contribution.

Chapter 2

Background

2.1 Multi-Agent Reinforcement Learning

Multi-Agent Reinforcement learning (MARL) allows for multiple autonomous agents to coexist in the same space. Formally, the setup of MARL consists of several *agents* that can perform specific actions, and the *environment* – a term that encompasses everything outside of the agents, that the agents can interact with. The specific formalisation of MARL that interests us in our work is called **Decentralized Partially Observable Markov Decision Problems** (in short: Dec-POMDPs) [20].

A Dec-POMDP is a cooperative process defined as a tuple $(\mathcal{N}, \mathcal{S}, \mathcal{A}, \mathcal{O}, \Omega, \mathcal{P}, \mathcal{R})$, where $\mathcal{N} = \{1 \dots N\}$ is the set of agents participating in the process. Agents interact with the environment in discrete timesteps $t \in \mathbf{N}$. In each timestep t the environment has some true active state $s_t \in \mathcal{S}$, and each agent i receives an observation $o_t^i \sim \Omega(i, s_t)$, $o_t^i \in \mathcal{O}$. Each agent i then selects an action $a_t^i \in \mathcal{A}$. After each timestep t the agents receive a shared reward $\mathcal{R}(a_t^1, a_t^2, \dots, a_t^N, s_t) = r_{t+1} \in \mathbf{R}$, and the environment enters the next state $s_{t+1} \sim \mathcal{P}(a_t^1, a_t^2, \dots, a_t^N, s_t)$, $s_{t+1} \in \mathcal{S}$. In figure 2.1 we include a diagram visualizing this process.

In MARL, and in reinforcement learning, the agents' goal at each point in time t is to maximize the discounted cumulative reward $\sum_{i=0}^{\infty} \gamma^i r_{t+i}$, where γ is a discount factor. When $\gamma = 1$, the discounted reward is equal to the actual reward $\sum_t r_t$ – we will report this sum when presenting evaluation results.

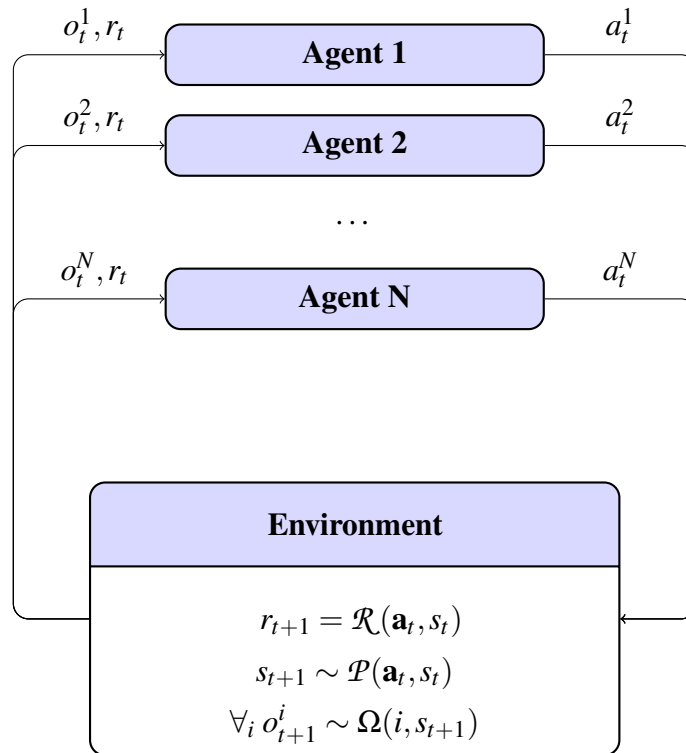


Figure 2.1: An illustration of the Dec-POMDP loop, adapted from a similar, single-agent figure in chapter 3 of the book by Sutton and Barto [18]. We are reusing this figure from our proposal [19] to help readers more easily understand the loop.

2.2 MARL algorithms

In our work, we wish to run a benchmark of several MARL algorithms using an environment we will implement. We will use the list of algorithms compiled in the benchmark paper by Papoudakis et al. [4], and reuse their implementations available in the ePyMARL framework for evaluation. In this subsection, we go briefly over the algorithms used and the main category split between them.

2.2.1 Independent Learning

In the **Independent Learning** approach, each agent treats each other agent as an element of the environment when learning. The algorithms falling into this category that we wish to run, following the benchmark paper, are: Independent Q-learning (**IQL**), introduced by Tan [21], which is an independent multi-agent version of Q-learning [22], Independent synchronous Advantage Actor-Critic (**IA2C**), which is an independent multi-agent version of the Advantage Actor-critic algorithm introduced by Mnih et al.

[23], and Independent Proximal Policy Optimisation (**IPPO**), which is an independent multi-agent variant of the PPO algorithm introduced by Schulman et al. [24].

2.2.2 Centralized Training, Decentralized Execution

This category of algorithms allows the training algorithm full knowledge of all of the agents' observations, as well as the true environment state, during training time. This allows for more informed decisions about updating the parameters of the agents. The agents are still, however, completely independent during the environment's execution.

One subcategory of algorithms following this approach are centralised policy gradient methods, in which a centralized "critic" network decides about the value of each state, while each agent acts upon that value function in a decentralized manner. This category includes Multi-Agent Deep Deterministic Policy Gradient (**MADDPG**), introduced by Lowe et al. [25], Counterfactual Multi-Agent (**COMA**), introduced by Foerster et al. [26], Multi-Agent Asynchronous Advantage Actor-Critic (**MAA2C**), a simple algorithm used as a baseline, explained further by Papoudakis et al. [4], and Multi-Agent Proximal Policy Optimisation (**MAPPO**), introduced by Yu et al. [15], an extension of the IPPO algorithm mentioned above.

Another subcategory are value decomposition algorithms, which aim to achieve a single joint value function by mixing the individual agents' value functions. This category includes Value Decomposition Networks (**VDN**), originally introduced by Sunehag et al. [27], and **QMIX**, originally introduced by Rashid et al. [14].

2.3 SMAC

2.3.1 Overview

In this section we go over the SMAC environment to the extent that is relevant to our project. Our project only replicates SMAC in its default configuration, omitting any optional parameters, so we omit those here as well.

In SMAC, each agent controls one unit (we refer to units controlled by agents as *allied units* throughout this paper) and is tasked with defeating a group of enemy units controlled by Starcraft II's built-in AI opponent. In Starcraft, a single combat unit is a single soldier in an army comprised of any number of units. SMAC defines several combat scenarios, differing in army compositions or terrain layout, and by result, in difficulty.




Figure 2.2: Visualizations of the SMAC environment. (left) The `corridor` scenario (right) The `MMM2` scenario.

Units are divided into several types with different attributes (health, attack, etc.) – but note that the SMAC environment only differentiates between unit types in the state/observation vectors if there is more than 1 unit type within a single team in the scenario – we will refer to this as the scenario “distinguishing unit types”.

Each unit uses one of two combat styles, determined by its type – either damage-dealing or healing. The only way to defeat a unit is to bring its health points down to 0 by attacking it, which is the job of damage-dealing units. Healing units can counteract that somewhat by regenerating their allies’ health points.

Some unit types, once hit, do not regenerate health in any way. Other units have innate health regeneration that is always active. Yet another group of unit types has special shields on top of their health points that regenerate after a period of not taking damage, and the shields have to be brought down to 0 before the units’ health can be hit. Due to limitations of Starcraft II, either all units in a team possess shields, or none of them.

A unit’s type also defines its attack range – some units (called *melee* units) need to get up close to an enemy in order to attack, while others can fire from a distance. Additionally, each unit has to wait a certain amount of time after attacking before it attacks again – this amount of time is called the unit’s *cooldown*.

2.3.2 Actions

Each agent has access to several actions, which may or may not be available at any given timestep – the environment makes available a method to get currently available actions for each agent. If an unavailable action is chosen by the agent, SMAC returns an error and ceases execution. The possible actions are:

`no-op`: Has no effect, this action is only available to dead units. All other actions are

only available to living units, so all dead units must always pick this action.

`stop`: Orders the unit to stop in its place and do nothing. The action is always available to alive units.

`moveN`, `moveE`, `moveS`, and `moveW`: Orders the unit to move in the chosen cardinal direction (north, east, south, or west). Each action `move{X}` is available to alive units only if a point near to the unit's current position in the direction `X` is within the bounds of the map and is not on rough (i.e. non-walkable) terrain.

`target1`, `target2`, ...: Orders the unit to target the unit with the specified team-specific ID – for damage-dealing units, this refers to targeting enemy units to attack, for healing units, to targeting allies to heal. SMAC defines a constant targeting range for agents – each action `target{i}` is available for alive units only if the target unit `i` is alive and within this targeting range.

Note that because the number of these actions varies by number of units, the total number of actions $|\mathcal{A}|$ varies in SMAC per specific combat scenario.

2.3.3 State

The true state of the environment is a vector divided into several sections, in this order:

- For each ally unit in order of their IDs: its current health, its current cooldown, its X and Y coordinates relative to the center of the map, its current shields (only if allies have shields), and a one-hot vector representing its unit type (only if scenario distinguishes unit types).
- For each enemy unit in order of their IDs: its current health, its X and Y coordinates relative to the center of the map, its current shields (only if enemies have shields), and a one-hot vector representing its unit type (only if the scenario distinguishes unit types).
- For each agent in order of their IDs: a one-hot vector representing the action taken by them in the previous timestep.

Note that all features within the state vector are normalized to be between zero and one – for example, the health value is divided by the maximum health value of the given unit.

2.3.4 Observations

Each agent, in each time-step, receives an observation representing what is visible to the agent within the environment. If the agent's unit is dead, the agent simply receives a vector of zeroes. SMAC also defines a constant sight range for agents – if a unit is dead, or is further from the agents' own unit than this sight range, any information in the observation vector about this unit is completely zeroed out. The observation vector includes, in this order:

- For each cardinal direction, whether movement is possible in that direction (if it is not, it could be because of rough terrain or because the unit is at the edge of the map).
- For each enemy unit in order of their IDs, if it is alive: whether the agent's unit can attack it (if it can't, the enemy is outside the agent attack range mentioned in the previous subsection), its distance to the agent's unit, its X and Y coordinates relative to the agent's unit, its health, its shields (only if enemies have shields), and a one-hot vector representing its unit type (only if the scenario distinguishes unit types).
- For each ally unit in order of their IDs, if it is alive: a literal 1 (to distinguish from units that are dead or too far), its distance to the agent's unit, its X and Y coordinates relative to the agent's unit, its health, its shields (only if allies have shields), and a one-hot vector representing its unit type (only if the scenario distinguishes unit types).
- For the agent's own unit: its health, its shields (only if allies have shields), and a one-hot vector representing its unit type (only if the scenario distinguishes unit types).

Note that, similarly to the state features, all observation features are normalized to be between zero and one.

2.3.5 Rewards

After each time-step all agents receive a shared reward equal to the sum of health points and shield points removed from enemies in that timestep. A small bonus of 10 is added for each eliminated enemy unit and a bigger bonus of 200 is added for winning the scenario. The reward is normalized by dividing it by the sum of all health and shield

points of enemy units and any possible bonuses, and multiplying by 20 – thus, the possible rewards should be between zero and twenty. Do note, however, that the actual cumulative reward received by the agents by the end of an episode might exceed 20 due to health and shield regeneration.

2.4 Optimal Reciprocal Collision Avoidance

2.4.1 Overview

In Starcraft II, units move around the battlefield populated by other units and avoid collision by stepping to the side if they would get in the way of another unit. This makes the battlefield feel more realistic and physical, and allows for some advanced strategies like body-blocking or surrounding (in essence, limiting other units' movement by positioning oneself strategically). We felt it was crucial to reproduce this behaviour in our environment. However, because Starcraft II is a proprietary, closed-source piece of software, we cannot use the exact algorithms used in SMAC. To fill this gap we chose the Optimal Reciprocal Collision Avoidance (ORCA) algorithm by Berg et al. [28].

The ORCA algorithm fills several criteria desirable for our environment. Much like in SC2, each unit is assumed to be a circle with a specific radius. Units can avoid other units, and they can also avoid static polygonal obstacles. On top of that, the units can not only avoid collisions, but also move towards their own goal location at the same time.

2.4.2 ORCA Methodology

The algorithm is parametrized by a time horizon τ . In each run of this algorithm, each unit A computes a set of half-planes (which we call *ORCA half-planes*) in 2D velocity space, each half-plane being the set of velocities safe to choose to avoid collision with some other unit or obstacle B for at least τ time. In this subsection, we briefly go over what the algorithm does to compute these half-planes.

Each unit A considers all units and obstacles in its immediate neighbourhood (i.e. within some radius r around it). For each neighbour B , the unit calculates the **velocity obstacle** induced by the neighbour on it – that is, the set of positions in velocity space that would make the unit collide (i.e. get within the distance of $r_A + r_B$, where r_A and r_B are the units' own radii) with that neighbour within τ time.

Let \mathbf{u} be the shortest vector from the relative velocity $\mathbf{v}_A - \mathbf{v}_B$ to the velocity obstacle's boundary – in other words, the smallest amount of change to the relative velocity required to prevent a collision within τ time. Then, the slope of the line (called the *ORCA line*) defining the ORCA half-plane for that neighbour is given by the slope of the outward normal of the velocity obstacle boundary at point $\mathbf{v}_A - \mathbf{v}_B + \mathbf{u}$. The line is anchored in a point given by $\mathbf{v}_A + \frac{1}{2}\mathbf{u}$. This gives the unit a half-plane of possible velocities that will avoid collision with the neighbour – note that the adjustment by $\frac{1}{2}\mathbf{u}$ is because the unit assumes the neighbour is following the same algorithm and will adjust by $-\frac{1}{2}\mathbf{u}$ (since from the neighbour's perspective, everything is mirrored – hence the negation) – the adjustment is not halved for obstacle neighbours, only for unit neighbours.

Then, given all half-planes induced by neighbours, the algorithm solves a linear programming problem for each unit, to find a velocity that avoids all neighbours and is the closest to the unit's desired velocity. If avoiding collisions completely is not possible, the algorithm solves a different linear programming problem that minimizes the distance the unit crosses behind the ORCA lines.

We go into more detail about how we use this algorithm in SMACLite in chapter 3. Note that this algorithm is not equivalent to a pathfinding algorithm – when faced with a wall, the units will often stop in front of it, they will not look for very far a way to go around it.

Chapter 3

Project Methodology

3.1 Environment implementation method

While in the previous chapters we described technology that was already available before SMAClite, and that we reused, in this chapter we focus on our contribution and how we approached implementing SMAClite itself. Do note that because of the closed-source nature of Starcraft II, we did not have access to any implementation details of it, and came up with the algorithms contained in this chapter via trial and error experiments, with the help of general information available on the Starcraft II Liquipedia [29].

We decided to implement the environment in the Python programming language [30] due to it being widely known in the machine learning community, where it is by far the most popular one. Most of the computations within the environment are performed using the Numpy library [31], and the rendering of the environment is handled by a script written by us using the Pygame [32] library. Wherever applicable, we used the default arguments of the SMAC environment, and omitted any optional ones. The SMAClite environment uses the well-established OpenAI Gym framework [33] for creating reinforcement learning environments.

At all steps of the implementation we made sure the action, state, observation, and reward APIs are exactly aligned with SC2/SMAC. This also applies for individual unit attributes, for which we consulted Liquipedia [29]. This allowed us to conduct transfer learning experiments, such as the one we describe in section 4.3.

Even though the grid in Starcraft II allows half-square triangles, we chose for simplicity to only allow a grid of squares as the terrain for SMAClite. We found that this simplifying assumption does not detract from the environment’s difficulty. Internally,

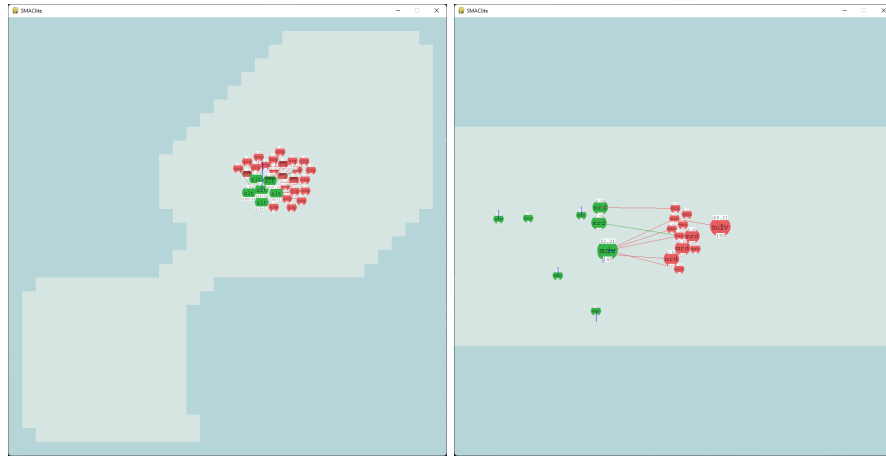


Figure 3.1: Visualizations of the SMAClite environment. (left) The `corridor` scenario (right) The `MMM2` scenario.

we collapse adjacent squares into rectangles to lower the total number of obstacles for performance’s sake. When defining the units’ velocities or sizes, the base distance measurement unit is the side length of a single square in the terrain grid. This is consistent with SC2, and all scenarios available in both SMAC and SMAClite use a map size of 32 by 32 units.

To further improve environment performance, at many points in time the units use K-D trees as available in the Scikit-learn library [34] to find their neighbours (e.g. finding units within sight range when generating observation vectors), as opposed to iterating over the entire unit list. Since K-D trees only support querying in a circular area and obstacles are always rectangles, when looking for obstacle neighbours, the units query an R-tree from the Python package `rtree` [35].

3.2 Base framework

SMAClite, much like SMAC before it, is defined mostly by the various combat scenarios available, as well as the units participating in those scenarios. As part of the SMAClite environment we contribute a framework capable of reading both custom scenarios and units from JSON files – this means expertise in the Starcraft II map editor is no longer required to create new or modified challenges using the environment. This also means that with SMAClite it is easy to tell the difference between two different unit types – to compare between ”zergling” and ”marine”, all one needs to do is look at their respective JSON definitions, and see what the differences are.

All of the standard scenarios and units shipped with the environment are written using this framework. In this section, we go over the details of what attributes are supported for both scenarios and units. Full examples of JSON files compatible with the framework can be found in Appendices A.1 and A.2

3.2.1 Scenario definition

Our framework accepts scenario JSON files containing a single JSON object. Each scenario should have a name to identify it, and should specify the numbers of allied and enemy units with the `num_allied_units` and `num_enemy_units` parameters. The units in the scenario should be listed with the `groups` parameter, each group being a JSON object with `x` and `y` parameters specifying the group's center, a `faction` parameter (ALLY or ENEMY) to specify which team the units are on, and a `units` parameter with an object of unit types together with their counts. Each group will initially be laid out in the shape of a square around their specified location. Note that these groups have no impact beyond unit positioning – once initial unit placement is complete, neither the agents nor the units have any information about what group they came from.

The framework supports two ways of specifying unit types. In order to use a standard unit type, its uppercase name should be used (e.g. ZERGLING – for a full list of supported standard unit types, see appendix A.3). The other way to specify a unit type is to provide a path to a JSON file with its specification. The framework also supports an optional `custom_unit_path` parameter – if specified, this path will be prepended to all unit types specified as a path. Note that if the `.json` extension is missing the framework reattaches it automatically, so if the custom unit path is `path/to` and the unit type specification is `type`, the framework will look for the file `custom/unit/type.json` file. The framework also requires an `attack_point`, which is typically near to the initial positions of the allied units – this point is where the enemy units will be marching towards throughout the scenario (see section 3.6 for details).

The framework also supports two ways of defining terrain – one way is to pick a standard terrain present by supplying the `terrain_preset` parameter with its uppercase name (e.g. CORRIDOR, see appendix A.4 for full list). Terrain can also be provided in the scenario definition file itself using the `terrain` argument, which should consist of a list of strings forming a rectangular 2D array. The framework supports two types of terrain: `_` for walkable, and `X` for non-walkable. The framework also always requires a `width` and `height` to be specified for the scenario, which should match the terrain

dimensions – note that all scenarios adapted from SMAC have the dimensions of 32 by 32.

Finally, the framework requires some general details about units in the scenario. The parameters `ally_has_shields` and `enemy_has_shields` specify which teams, if any, should have shields on them, and the parameters `num_unit_types` and `unit_type_ids` specify what IDs the agents will receive in observations for the various unit types. The former should be a number (note that scenarios adapted from SMAC use 0 for all scenarios where both teams are homogenous, i.e. only have one unit type each), and the latter should be a map from unit type specifications (same as above) to numbers from 0 to `num_unit_types` minus one. The length of the map must be equal to `num_unit_types`.

We ship several standard scenarios with this environment as separate OpenAI Gym environments (e.g. `smaclite/2s3z-v0` – see appendix A.5 for the full list of 12 scenarios we adapted from SMAC). To use a custom scenario file, one should use the `smaclite/custom-v0` Gym environment, and provide a path to the scenario file via the `map_file` parameter.

3.2.2 Unit definition

Each unit is assigned a specific **type**, with a set of different attributes impacting the environment mechanics. Our framework supports several attributes in the JSON files defining the various unit types, in order to allow for their easy customization.

Firstly, the framework supports several attributes defining the units' resources, including their maximum health and health regeneration via the `hp` and `hp_regen` attributes, respectively. Their shields, if any, via the `shield` attribute, and their maximum and initial energy via the `energy` and `initial_energy` attributes. The units' size can be specified by providing their diameter in the `size` attribute, and their speed (in distance per second) via the `speed` parameter.

Then, the framework supports several attributes defining the units' combat abilities. Their `combat_type` (one of `DAMAGE` or `HEALING`) defines their main role on the battlefield, and their `damage` defines how strong of a hitter they are, while their `armor` defines their defensive capabilities. Each unit has a specific `attack_range`, which defines how close (measured boundary to boundary) the unit has to get to a target in order to attack or heal it¹. Note that while a numeric value is usually expected for this attribute, the special value of `MELEE` is also accepted to signify that the unit should have the standard

¹Note that this is different from the agent targeting range mentioned in section 2.3 – that range only affects the agents' observations, while this one actually governs when an attack can happen.

melee range, i.e. only attack from up close. Some units can deal damage multiple times per attack – this can be achieved using the `attacks` attribute (e.g. 2 for attacking twice at once), and after attacking, each unit has to wait their specific `cooldown` before attacking again. Each unit type should also define a `minimum_scan_range`, which will govern how far the units will look when searching for targets.

In SMACLite, each unit resides in a specific plane, with three currently supported: `GROUND`, `AIR`, and `COLOSSUS`. When moving, the units only avoid collisions with units in the same plane as them, and only ground units are affected by static obstacles. All of the units in SMACLite can only target units which reside in the planes listed in their `valid_targets` – for example, if `AIR` is not in this list for some unit, then it can never attack airborne units. But note that, regardless of their `valid_targets`, all units can target units in the `COLOSSUS` plane.

In order to support making certain unit types stronger against specific other types, the framework also supports a system of unit attributes (e.g. `BIOLOGICAL`, see appendix A.6 for a full list of available attributes) and attribute-relative bonuses, defined as a map from attribute to bonus value. For example, if a unit type has a bonus of 20 against `ARMORED` units, it will deal 20 bonus damage with each attack against units with that attribute.

Finally, the framework supports several different attack types for units, defined by the `targeter` and `targeter_kwargs` attributes. The most common attack type is `STANDARD`, which simply has the unit attacking one target at a time. The other attack types are only used by one standard unit each, but can easily be reused for custom unit types. The `KAMIKAZE` attack type has the unit explode when attacking, dealing damage in a circle with a specified `radius` around itself and dying the process. The `LASER_BEAM` type fires a laser in a line perpendicular to the line between the attacker and their target – more specifically, the laser line is a rectangle with a specified `width` and `height`. Lastly, the `HEAL` attack type is used by healer units.

3.3 Unit command types

At any given point in time, each unit in the environment is executing one of several types of commands. SMACLite supports five different command types, with the first four being exactly equivalent to the four action types available to the agents, as described in section 2.3.

We introduce one more command type that is unavailable to agents but is key to the

AI opponent's behavior, called `attack_move` – this command orders the units to march toward a specified location, attacking any units encountered along the way, and then guard the location once it is reached. We based our implementation of attack-moving on the "Automatic targeting" article on the Starcraft II Liquipedia [29], but made a few judgement calls based on what yielded the most desirable behaviour, wherever their information was unclear or unavailable.

3.4 Game loop

Each environment step starts with the agents' units being assigned commands corresponding to the actions chosen by the agents. Once that happens, we simulate eight game steps, and the reward returned from the environment step is the sum of the rewards earned within these game steps. Each game step is considered to last $\frac{1}{16}$ th of a second for the purpose of calculating velocity, cooldowns, etc. – that does not mean this is its actual duration, as in fact our environment can run much faster than real-time (see 4.4.1 for details).

Each game step consists of several phases. These phases were developed by us and we have no information about whether the game steps inside Starcraft II follow any sort of similar order – this structure was simply what yielded the most reasonable game ruleset that resembles Starcraft II. Before the next phase can begin, all units must execute the logic for the previous phase. This is necessary to ensure the units' perceptions of other units (note: unrelated to the agents' observations) remain consistent throughout the execution of the game step.

3.4.1 Target clean-up

First of all, each unit might lose the target it was attacking or healing in the previous game-step, according to command-specific logic. It is necessary for all units to execute this logic before proceeding with declaring their preferred velocity for this game step because other units might access the target information when computing their own preferred velocity.

Units with the `noop`, `stop`, and `move` commands immediately lose any target they had, since these are non-combat commands. Units with the `target` command acquire the target dictated by their command, or retain it they are already targeting it. With the `attack_move` command, the units consider several factors. They always lose a dead

target and they never lose a target who attacked them in the last game step. Otherwise, they lose their target if it is outside of their attack range.

3.4.2 Velocity preparation

In this phase all units declare their preferred velocity for this game step. It is necessary that they all do this before proceeding with velocity adjustment via the ORCA algorithm, because units can perceive each other's preferred velocity and make collision avoidance decisions based on that information.

In the case of `noop` and `stop` commands, the preferred velocity of the units is always $\mathbf{0}$. In the case of the `move` command, if the unit's maximum velocity is v_{max} , its current position is \mathbf{x} , and the position where it wants to move is \mathbf{y} , where $\mathbf{x} \neq \mathbf{y}$, then the unit always declares that its preferred velocity is $(\mathbf{y} - \mathbf{x}) \frac{v_{max}}{\|\mathbf{y} - \mathbf{x}\|}$. If $\mathbf{x} = \mathbf{y}$, then the preferred velocity is instead $\mathbf{0}$. Put simply, the unit always wishes to move in a straight line towards its destination – this is different to Starcraft II's engine which has a built-in pathfinding algorithm based on the A* algorithm [36] – but we decided to omit it, since it only matters in one of the many scenarios offered by SMAC (`2c_vs_64zg`).

If the unit's command is `target`, the unit considers its distance from the target unit. Let the unit's attack range be d_{max} , its radius r_A , and its target radius r_B , and let the distance between it and the target be defined as $d(A, B)$. Then, in the case where $d(A, B) > r_A + d_{max} + r_B$, the unit declares it is moving in a straight line towards its target, and proceeds as with a `move` command. Otherwise, it declares it is attacking or healing, and its preferred velocity is $\mathbf{0}$.

Finally, if the unit's command is `attack_move`, the unit first finds the valid targets within its scan range. It is important to note that if the unit is a damage-dealer, it considers any enemy healers *priority targets*, and all other enemy units as non-priority targets – if the unit still has a target after the target clean-up phase, the only way it will switch targets at this point is if its current target is not a priority target, and there is a priority target within its scan range. If it does not have a target, or needs to switch, it picks the closest target with the highest available priority among the valid targets. The process is slightly different with healers, who consider valid targets any non-healer units in their team who are below full health, or who are attacking another unit. The healer unit picks as its target the lowest-health unit among its valid targets. After the target selection process finishes, the unit proceeds as with a `target` command if it has a target, and as with a `move` command toward its attack-move destination position if

it does not. Note that healers will adjust their maximum velocity for any given game step to match the slowest allied unit within their scan range – we implemented this behaviour in order to stop them getting in front of their army if they are the fastest units.

There are a few intentional changes from SC2 in this phase – therein, it is not true that healers can target any non-healers, and can never target other healers or themselves. This is only true with the limited set of units present in SMAC, and we felt it makes for a nice simplification of the SC2 ruleset that does not change anything in SMAC. It is also not true that damage-dealers consider any healers as priority targets in SC2 – they only do in SMAC due to the modifications made to the map files by its authors – again, we felt this is a nice simplification aligned with SMAC, but not with SC2.

3.4.3 Velocity adjustment

Each unit A uses the ORCA algorithm to determine its actual velocity that avoids collisions, taking into account its own preferred velocity, as well as its neighbour units and obstacles. For the purpose of collision avoidance, the unit considers its neighbours all units within the radius of $(r_A + r_{max})\tau$ of itself, where r_{max} is the maximum radius of any unit in the scenario and τ is the time horizon; and all obstacles within the radius of $r_A + \tau v_{A,max}$ of itself. In SMAClite, we always use $\tau = 1$, i.e. the units want to guarantee avoiding collisions for 1 second, or 16 game steps. The ORCA algorithm returns for each unit its actual velocity, given its preferred velocity. For details on our implementation of the ORCA algorithm, refer to section 3.5.

3.4.4 Game step execution

This is the final phase of the game step, and because order matters here (a unit might be eliminated before it gets to execute its command), the units execute their respective logic for this phase *in random order*, with the order being re-randomized in each game step. Note that we do not have any data about whether Starcraft II also randomizes game step execution order, this simply felt like a decent compromise between fairness and simplicity.

Each unit, regardless of its command type, begins this phase by performing several standard updates. It updates its position according to its actual velocity computed in the previous phase (i.e. $\mathbf{x}_{new} = \mathbf{x} + \frac{\mathbf{v}_{actual}}{16}$, since velocity is always defined in distance per second), it reduces its cooldown if applicable (i.e. if it is greater than 0, it gets reduced by $\frac{1}{16}$ th of a second), and it regenerates health, energy, and/or shields, if applicable. It

then proceeds with command-specific logic.

The only case where any command-specific logic is performed at this point is if the unit's command is `target`, or `attack_move` with a non-empty target, and the unit declared during the velocity preparation phase that it is attacking or healing, and its current cooldown is 0 – in all other cases the unit's game step logic ends here. In the single actionable case, it attacks or heals its target.

Damage from attacks is dealt first to the enemy shields, then to the enemy health, with any damage dealt to the target's health being reduced by its armor. If the unit's `attacks` attribute is greater than 1, then it deals damage multiple times in a row during this step. If the unit did attack, its cooldown is then set to its maximum value, defined by the unit's type. Healing units heal their target at a rate of 9 health per second, spending $\frac{1}{3}$ of an energy point per health healed.

Do note that units might attack or heal units that are technically outside of their range at this point, because attack/heal declaration happened during the velocity preparation phase, and the target might have moved away slightly since then. This is intentional and is meant to prevent endless chases when the units' velocities are similar. Again, we do not have data on whether the Starcraft II engine features any similar simplifications, but this rule works well for SMACLite.

In addition, there is an intentional change from SC2 in here as well, because for simplicity we do not simulate attack animations or attack projectiles – every attack happens instantly in the timestep when the unit declares it is attacking. This causes some small discrepancies from SC2, but we believe the margin of error is acceptable.

3.5 Custom ORCA implementation

To facilitate installation simplicity, we rewrote the RVO2 library – the official implementation of ORCA in the C++ programming language from Berg et al. [28] – into Numpy, and we ship this module together with SMACLite. Because Starcraft II uses compiled (and, we assume, highly optimized) C/C++ code, and Python code can be quite slow compared to it, we ran into performance issues early in the development of SMACLite. To address this issue, we also make available an addon for SMACLite, which uses the original C++ RVO2 library verbatim, using Python bindings written in the Cython [37] extension. The Python bindings were originally made available by Stüvel [38], though we implemented several modifications to suit our use-case – these modifications are also present in the Numpy version of RVO2.

The original RVO2 library has no way to remove units one by one, nor remove all units at once. These are key features for us, as we need to adjust the collision avoidance unit list whenever a unit dies or whenever we restart the environment. We added both of these features into RVO2 – we assume dead units disappear from the battlefield as soon as they are eliminated, so we do not want other units avoiding collisions with them.

The second set of adjustments considers static units – the original ORCA algorithm assumes fully cooperative units who will go out of their way to make the passage easier for other units. Our version, on the other hand, assumes that if a unit A is static (i.e. $\|\mathbf{v}_A\| = 0$), it will *never* adjust its velocity. This makes it possible to surround other units and/or block their path, which is a valid strategy in Starcraft II. In the original implementation of ORCA, the blocking units would simply be ”pushed” away. Note that to make up for this, any moving units will adjust their velocity by \mathbf{u} instead of $\frac{1}{2}\mathbf{u}$ (see subsection 2.4.2 for definition of \mathbf{u}), when avoiding static units – this ensures moving units will not walk into static units.

Because the addon requires building and installing C++ files via CMake, which could be problematic on some systems, we chose to make this RVO2 fork available as an optional dependency, for users who are willing to go through a more difficult installation process in order to boost environment performance. We call this version `SMAClite_plus`. Note that, because of differences in finding neighbours, unit behaviours will not be exactly the same between the two versions, but should remain functionally indistinguishable. If this addon is in use, custom K-D trees implemented in C++ for RVO2 are used for finding unit and obstacle neighbours for collision avoidance purposes, instead of Scikit-learn K-D trees or `rtree` R-trees. After it is installed, the addon can be enabled by setting the parameter `use_cpp_rvo2` to `True` when initialising the environment. Any attempt to set this argument to `True` without the addon installed will result in an error.

3.6 Opponent AI behaviour

The authors of SMAC [5] claim the opponent team is controlled by Starcraft II’s built-in AI on the *very difficult* level. Because the enemy units’ behaviours seemed very simple, we had our doubts about their strategic ability when watching combat inside the SMAC environment. For this reason, we performed two experiments:

1. First, we toggled the AI level inside SMAC across the various difficulty levels available (including the ”cheating” levels), while keeping the random number

generator seed constant. We then put the opponent AI against agents who pick randomly from the available actions. We saw no difference at all in the rewards obtained by the random agents between difficulty levels – they were exactly equal.

2. Then, going one step further, we **removed** the AI opponent from the game, making the enemy units not controlled by any player. This **did not change the rewards either**.

We performed these experiments on the `MMM2`, `2c_vs_64zg`, and `corridor` SMAC scenarios.

Based on these results, we are reasonably certain the AI never issues any orders to the enemy units in SMAC. The only order given to them is inside a script in each SMAC map file, which tells them to `attack-move` towards a specific point, usually where allied units initially appear. Following these results, we did not implement any custom opponent AI. All we do is, upon map initialization, set the enemy units' command to `attack_move` towards the `attack_point` specified in the map scenario file. The enemy units' command never changes throughout the encounter.

Chapter 4

Experiments

In this section we describe the results of various experiments we performed in our environment. In all of the experiments, we used the scenarios listed in table 4.1. More specifically, we use all of the scenarios used by Papoudakis et al. [4] in their paper, with the addition of `bane_vs_bane`, which we included to feature a wider selection of unit types.

4.1 Agent learning curves

We trained all of the algorithms listed in section 2.2 on the selected scenarios. Each training was run for 4 million steps, and each training was repeated using 5 different random number generator seeds. All of the trainings were performed on the `SMAClite_plus` version of the environment, and were performed solely using CPUs, each training using a single CPU core. The trainings were run on various nodes on a cloud cluster, most of them using Intel(R) Xeon(R) Gold 6138 CPUs @ 2.00GHz, and all others used CPUs of comparable computational capacity. Due to technical constraints, all trainings were under a strict 48-hour time limit. For all of the algorithms, we used hyperparameters listed as the best found for SMAC in the benchmark paper [4]. For the episodic algorithms we used a buffer size of 5000, and for the parallel algorithms we used a buffer size of 10. We used the ePyMAREL framework to run all of the trainings.

Some minor caveats need to be mentioned with regards to the learning curves. Firstly, the `bane_vs_bane` scenario is by far the slowest (as evidenced by subsection 4.4.1), so we reduced the required timestep to 1 million for this scenario only. In addition, we found the MADDPG algorithm extremely slow during training, and so it never reached the required timesteps in two scenarios (`MMM2` and `bane_vs_bane`) under

the time limit, so we omit those two curves in our figure.

The resulting graphs of mean reward over training time (also known as learning curves) can be found in figure 4.1. In the remainder of this section, we compare the learning curves to those reported for SMAC in the benchmark paper.

First of all, many algorithms easily solve the `2s_vs_1sc` scenario in SMAC, quickly reaching the maximum reward of 20, while in SMAClite the best performer (QMIX) only reaches a mean reward of about 15. We theorize that this is due to the fact that SMAClite does not simulate attack animations, and originally in Starcraft II the spine crawler has a very long attack animation (0.238 seconds versus 0.1193 seconds for the stalkers), so in SC2 it is much easier to dodge the spine crawler's attack in the last possible moment. Though the difference is less noticeable there, a similar situation occurs in `3s_vs_5z`, and we believe it is for the same reason. Put simply, we believe *kiting* – alternating between running away and attacking without getting hit – is much more difficult in SMAClite than in SC2 because of the instant attacks.

Another major difference between the SMAC and SMAClite learning curves is in the `MMM2` environment – in the case of SMAClite, three algorithms (MAPPO, VDN and QMIX) seem to have mastered the environment, with maximal rewards almost reaching 25 (for details about how the algorithm achieves this see the next section), while in SMAC the best performing algorithms barely reach a reward of 17.5 (but note that they are the same three algorithms, which shows the scenarios are still somewhat equivalent).

Our hypothesis is that this is due to a potential bug in the SMAC environment code. When calculating rewards, SMAC simply subtracts enemy units' new health values from their old health values, so any healing done by the enemies results in a negative reward being incurred by the agents, even though they never did anything wrong. We believe that this, combined with the fact that rewards should be gained for dealing damage, could confuse the agents during training. Judging by the wording in the SMAC paper [5] we do not believe this is intentional, so the same behaviour is not present in SMAClite, where we explicitly define the reward as the sum of health points lost by the enemies due to attacks, and we do not penalize the agents for enemies healing.

Despite these few issues, we do observe that the overall shape of the training curves is comparable in all scenarios, some of them differ only in maximum reward reached. We believe that this is promising evidence pointing towards the environments being equivalent as far as learning is considered.

Name	Allied Units	Enemy Units	Description
2s_vs_1sc	2	1	Two stalkers – powerful but fragile ranged units – face off against one spine crawler – very strong unit with no movement capabilities. The stalkers need to abuse the spine crawler’s immobility to bring its health down.
3s5z	8	8	A symmetrical map with each team having three stalkers and five zealots – melee units that deal less damage but can take a lot of hits before dying.
MMM2	10	12	Each team has some marines – fragile ranged units that attack quickly, some marauders – more durable units that have stronger attacks, but can’t hit flying units, and one medic – a flying healer unit.
corridor	6	24	The allied team only has 6 zealots to hold off 24 zerglings – extremely fragile but very quick units with potential to overwhelm unprepared enemies with numbers and surround them.
3s_vs_5z	3	5	Three stalkers need to keep the quick zealots at bay while outnumbered, using their range to their advantage.
bane_vs_bane	24	24	Each team has some zerglings and some explosive kamikaze banelings, each of which can easily take out several zerglings with one explosion.

Table 4.1: Combat scenarios used in experiments described in this section.

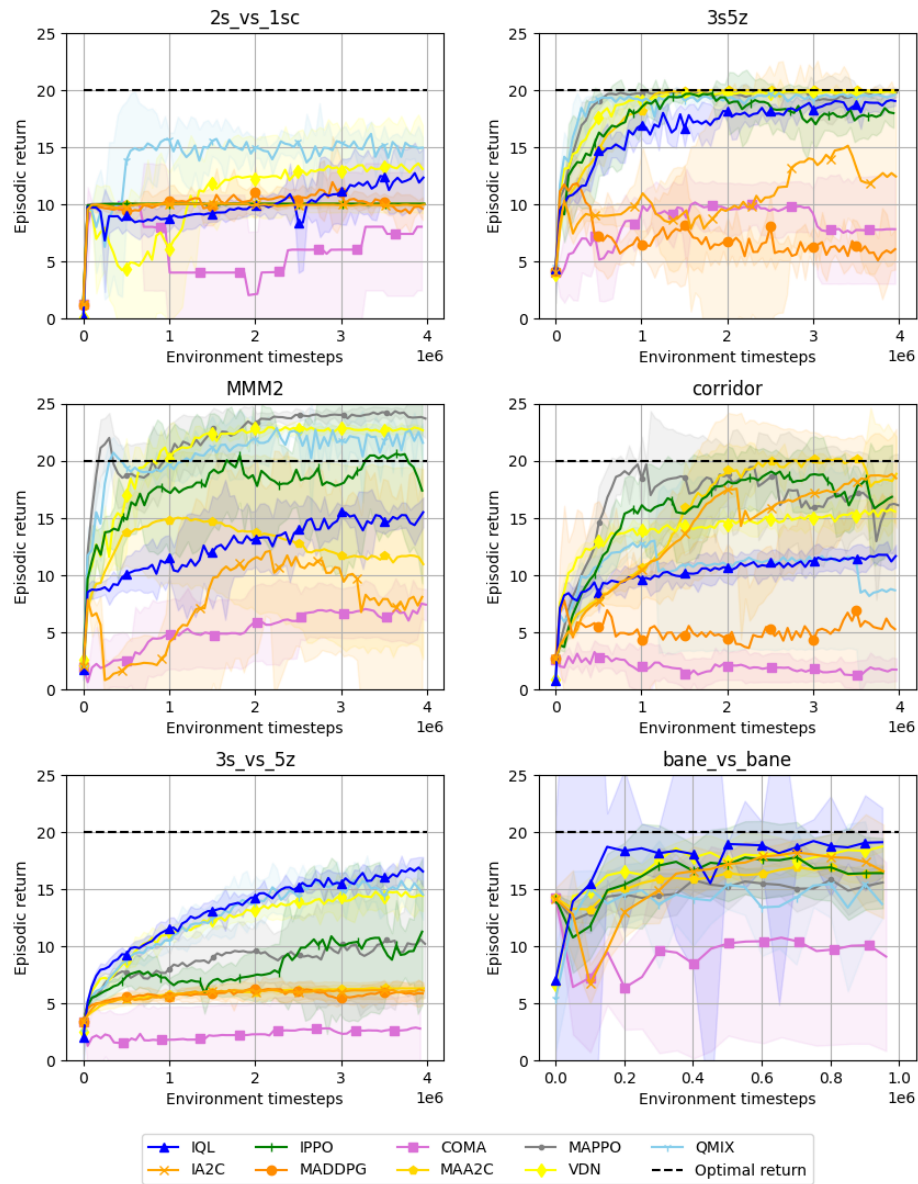


Figure 4.1: Test-time rewards received by agents trained using the different algorithms over time during training. The graphs show the mean value, as well as the 95% confidence interval, from 5 trainings differing by random number generator seed.

4.2 Behaviours observed in agents

In this section we select one model for each scenario, and we describe the strategies employed by the agents to achieve high reward values, in order to demonstrate to the readers the complexity of strategies required in SMAClite. Note that the below descriptions are simply our interpretations of the behaviours demonstrated by the agents in the environment, informed by our knowledge of the game gained both by playing Starcraft II and working on SMAClite.

When making our selection, we wanted to showcase each of the algorithms, and each of them in a scenario where it did well – note that this means we might not use the best performer in all of the scenarios, but all of them are at least above average. However, because there are 9 algorithms and only 6 scenarios and because not all algorithms performed well, we decided to omit COMA, MADDPG, and IPPO, which were the overall weakest performers. In each case we used the latest available (i.e. highest amount of training timesteps) checkpoint of agent parameters.

`2s_vs_1sc` – QMIX – **Mean test reward: 16.22** – one of the stalkers is the "baiter" gets the spine crawler's attention and they both attack until the baiter's health drops down to a very low value, at which point it backs out and lets the spine crawler target the other stalker. They then both hit the spine crawler until they die, which sometimes results in a victory and sometimes in defeat depending on the random attack order.

`3s5z` – VDN – **Mean test reward: 20** – the stalkers and the zealots both focus on the enemy stalkers first, which deal more damage per second than the enemy zealots. The allied zealots move slightly north and the allied stalkers move slightly south, both to direct the attention of enemy stalkers (whose initial position is on the north side of the enemy army) to the allied zealots, who can take more hits, and also to allow allied stalkers plenty of room to maneuver while attacking. Also worth noting is the fact that the allied units are quite good at moving away right as their health drops to a dangerously low level, letting other nearby allies take the aggression.

`MMM2` – MAPPO – **Mean test reward: 24.63** – all allied damage-dealers focus on the enemy marauders first, because they are the heavier hitters, while the ally medivac hides behind others as soon as it gets low, since it is a priority target and would quickly be gunned down. What is interesting here is that the allied units

leave the enemy medivac alive and kill it *last*, abusing the fact that enemy units getting healed result in a higher reward due to the total damage dealt being higher – this is how the agents surpass the optimal return of 20 and reach values close to 25. They even stop hitting enemy damage-dealers and let them get healed, since the medivac cannot heal itself.

corridor – MAA2C – **Mean test reward: 20.25** – the zealots fan out and form a horizontal line against the zerglings, blocking them making it harder for the zerglings to surround them. This causes the zerglings to crowd around the front of the zealot line, reducing the total amount of damage the zealots take over time.

3s_vs_5z – IQL – **Mean test reward: 18.1** – the stalkers use the intended optimal strategy of kiting the zealots around the map. The lowest-health stalker always makes sure to stand behind the other two when attacking, in order to avoid dying. The stalkers do seem to get "lazy" with their kiting when the number of remaining zealots becomes low, probably because there is no penalty for dying, and it is easier to just stand and attack when the risk of death is low.

bane_vs_bane – IA2C – **Mean test reward: 19.14** – the zerglings run away to the west to avoid the enemy banelings' explosions, while the ally banelings charge forward and explode when they become surrounded by enemy units. The enemy zerglings and banelings quickly die to multiple explosions. If any enemies remain after the allied banelings' explosions, the allied zerglings come out of hiding and attack.

4.3 Cross-environment 0-shot performance

In order to verify whether our environment does in fact maintain the same features as SMAC, we performed 0-shot transfer learning experiments on each scenario (i.e. without any retraining on the environment we transfer to). We used the same models as in the previous experiment – to be specific, we first used their parameters from the first time the model was saved (at the very beginning of the training process), and then their parameters from the last time the model was saved (at the very end of the training process), and then compared the mean test-time rewards.

We note that the mean reward obtained has increased in *all* of the scenarios, with some exhibiting significant improvements such as doubling or tripling of the mean

Scenario	algorithm	early reward	late reward	late reward on SMAClite
2s_vs_1sc	QMIX	0	11.4	16.22
3s5z	VDN	3.09	8.96	20
MMM2	MAPPO	1.87	6.70	24.63
corridor	MAA2C	3.35	4.48	20.25
3s_vs_5z	IQL	3.13	9.33	18.1
bane_vs_bane	IA2C	18.97	19.84	19.14

Table 4.2: Mean test rewards achieved by agents trained on SMAClite when put inside the original SMAC environment, achieved using the parameters from the first time they were saved during training, and from the last time they were saved during training. Late reward on SMAClite also included for reference.

reward. A noteworthy example is the `bane_vs_bane` scenario, where even the early version of the agents got close to the optimal return of 20, but was still improved upon by the later version. Because all of the scenarios exhibited improved rewards upon training in SMAClite, we believe there is evidence to support the two environments requiring similar sets of skills, and that transfer learning from one to the other is a viable training strategy.

4.4 Environment Performance Benchmark

In this section, we consider the performance of the environments themselves, to confirm or deny whether SMAClite is indeed cheaper to run than SMAC. To obtain the data in this benchmark, we ran each scenario 20 times with agents picking randomly among the available actions. All of the below experiments were run on a computer with an AMD Ryzen 3700X CPU.

4.4.1 Environment speed

In table 4.3 one can find the time each environment took per timestep (excluding logic not belonging to the environment, like action selection). We notice that the pure Python code is indeed slower than the original SMAC environment, but the environment becomes much faster when using the C++ RVO2 addon.

Note that these timings correspond to running 1 environment step, which consists

Scenario	SMAC	SMAClite	Change Factor	SMAClite_plus	Change Factor
2s_vs_1sc	0.004	0.007	175%	0.003	75%
3s5z	0.013	0.018	138%	0.007	54%
MMM2	0.017	0.028	164%	0.010	59%
corridor	0.014	0.092	657%	0.010	71%
3s_vs_5z	0.006	0.013	216%	0.005	83%
bane_vs_bane	0.049	0.086	176%	0.024	49%

Table 4.3: Average seconds per timestep on the SMAC scenarios we used for our trainings. Data was obtained by running the scenario 20 times against random agents.

of 8 game steps in sequence, together with any environment-only logic (e.g. calculating rewards and determining observations). If we ever were to run experiments against human players (assuming SMAClite got some human control extension), we would want the environment to be capable of running in real time. Both in SC2 and SMAClite, each game step is considered to last $\frac{1}{16}$ seconds, and most human players play SC2 at the "faster" in-game speed, which corresponds to a 40% speed-up. Therefore, in order for the environments to run in real-time, they can afford to use $\frac{8}{16*1.4} \simeq 0.357$ seconds per environment step. Thus, all versions of the environment are more than capable of running all tested scenarios in real-time – the advantage of SMAClite_plus becomes the most obvious when running lengthy trainings which require many executions of the environment.

4.4.2 Environment memory usage

While running the experiments from the previous section, we also measured the amount of memory used by each environment. This did not vary a lot by scenario, and oscillated around 600 MB for the SMAC environment, and around 100 MB for the SMAClite/SMAClite_plus environments. Therefore, our lightweight version of the environment requires six times less memory than the original SMAC environment to run.

Chapter 5

Conclusions

5.1 Future Work

We believe the SMAclite engine and framework are both very much open to extensions. Thanks to the fact that the environment is no longer bound by the Starcraft II dependency, developers could introduce game mechanics unrelated to Starcraft II, and could also go beyond the game’s technical limitations. One could implement new unit types unseen in SC2 with new attack types or abilities, or easily create countless intricate puzzles for the agents to solve using the scenario framework.

It is also possible to treat SMAclite strictly as an extension of SMAC, and work towards making it as close to SMAC as possible, while maintaining the performance improvements it brings. This would likely mean focusing on transfer learning experiments such as the one in section 4.3, and eliminating various differences present currently between the environments. One example might be attack projectile/animation simulation, since that is one of the major differences between the two, and caused discrepancies in `2s_vs_1sc`, as well as in `3s_vs_5z`, during our experiments.

One could also work on improving the modified ORCA algorithm’s combat capabilities – as evidenced by the `corridor` scenario, the SC2 zerglings were much better at surrounding the zealot wall as performed by our MAA2C agents than the SMAclite zerglings. This probably stems from the fact that RVO2 is mostly a general-purpose collision avoidance algorithm, while SC2’s algorithm was handwritten for the best combat performance possible. Another feature idea we could use from SC2 is a pathfinding algorithm, since running in a straight line becomes a problem very soon when terrain becomes any more complicated than the terrain in our experimental scenarios.

5.2 Summary

In this paper, we presented a Multi-Agent Reinforcement Learning environment based on the Starcraft Multi-Agent Challenge [5]. We implemented it from the ground up using the Python programming language, mostly using the Numpy and Scikit-learn libraries. A major feature of the environment, other than complete decoupling from the heavy dependence on the Starcraft II game, is a framework capable of reading custom scenario and unit files defined in the widely used JSON format, allowing for easy interpretation and modification of the various scenarios and units present in the environment, as well as equally easy introduction of new ones.

We also introduce and make available two versions of the RVO2 library by Berg et al. [28] – one implemented using Numpy for maximum convenience, and one being a modification of the original C++ library optimized for performance. Both versions are compatible with our environment out of the box. Depending on priorities, each of them can be used interchangeably, and both offer real-time collision avoidance for all of the units in the environment.

We trained 9 different MARL algorithms in our environment using a modified version of the ePyMARL framework, which we also make available. We presented the learning curves of these algorithms on 6 different SMACLite scenarios, comparing these results to corresponding SMAC results in a benchmark paper, and noting that all of the learning curves exhibit similar properties to the ones from the original SMAC environment. We then offered verbal descriptions of the strategies employed by the MARL agents on each scenario to demonstrate the depth of strategy possible to achieve in SMACLite. We also performed a transfer learning experiment where we showed training agents on SMACLite brings improved performance in SMAC in all scenarios tested.

Following that, we ran environment performance experiments to verify that SMACLite is indeed a cheaper environment to run than SMAC. We observed that SMACLite with its C++ RVO2 variant was faster than SMAC on all scenarios tested, and both versions of the SMACLite environment were using six times less memory than SMAC on all scenarios tested.

We believe that the SMACLite environment is equally as challenging as SMAC, and can serve as equally good of a standard MARL benchmark. It is our hope that the improved accessibility and performance that SMACLite offers will be viewed as strong arguments for making it the benchmark of choice for MARL algorithm evaluation.

Bibliography

- [1] F. Vaussard, J. Fink, V. Bauwens, P. Réturnaz, D. Hamel, P. Dillenbourg, and F. Mondada, “Lessons learned from robotic vacuum cleaners entering the home ecosystem,” *Robotics Auton. Syst.*, vol. 62, no. 3, pp. 376–391, 2014. [Online]. Available: <https://doi.org/10.1016/j.robot.2013.09.014>
- [2] S. V. Albrecht, C. Brewitt, J. Wilhelm, B. Gyevnar, F. Eiras, M. Dobre, and S. Ramamoorthy, “Interpretable Goal-based Prediction and Planning for Autonomous Driving,” in *IEEE International Conference on Robotics and Automation, ICRA 2021, Xi’an, China, May 30 - June 5, 2021*. IEEE, 2021, pp. 1043–1049.
- [3] F. Christianos, L. Schäfer, and S. V. Albrecht, “Shared Experience Actor-Critic for Multi-Agent Reinforcement Learning,” in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M.-F. Balcan, and H.-T. Lin, Eds., 2020. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/hash/7967cc8e3ab559e68cc944c44b1cf3e8-Abstract.html>
- [4] G. Papoudakis, F. Christianos, L. Schäfer, and S. V. Albrecht, “Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks,” in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, J. Vanschoren and S.-K. Yeung, Eds., 2021. [Online]. Available: <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/a8baa56554f96369ab93e4f3bb068c22-Abstract-round1.html>
- [5] M. Samvelyan, T. Rashid, C. S. d. Witt, G. Farquhar, N. Nardelli, T. G. J. Rudner, C.-M. Hung, P. H. S. Torr, J. N. Foerster, and S. Whiteson, “The StarCraft Multi-Agent Challenge,” in *Proceedings of the 18th International Conference on*

- Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, E. Elkind, M. Veloso, N. Agmon, and M. E. Taylor, Eds. International Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 2186–2188. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3332052>
- [6] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vechnevets, M. Yeo, A. Makhzani, H. Küttler, J. P. Agapiou, J. Schrittwieser, J. Quan, S. Gaffney, S. Petersen, K. Simonyan, T. Schaul, H. v. Hasselt, D. Silver, T. P. Lillicrap, K. Calderone, P. Keet, A. Brunasso, D. Lawrence, A. Ekermo, J. Repp, and R. Tsing, “StarCraft II: A New Challenge for Reinforcement Learning,” *CoRR*, vol. abs/1708.04782, 2017, arXiv: 1708.04782. [Online]. Available: <http://arxiv.org/abs/1708.04782>
- [7] M. Plappert, M. Andrychowicz, A. Ray, B. McGrew, B. Baker, G. Powell, J. Schneider, J. Tobin, M. Chociej, P. Welinder, V. Kumar, and W. Zaremba, “Multi-Goal Reinforcement Learning: Challenging Robotics Environments and Request for Research,” *CoRR*, vol. abs/1802.09464, 2018, arXiv: 1802.09464. [Online]. Available: <http://arxiv.org/abs/1802.09464>
- [8] E. Todorov, T. Erez, and Y. Tassa, “MuJoCo: A physics engine for model-based control,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012*. IEEE, 2012, pp. 5026–5033. [Online]. Available: <https://doi.org/10.1109/IROS.2012.6386109>
- [9] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling, “The Arcade Learning Environment: An Evaluation Platform for General Agents,” *J. Artif. Intell. Res.*, vol. 47, pp. 253–279, 2013. [Online]. Available: <https://doi.org/10.1613/jair.3912>
- [10] I. Mordatch and P. Abbeel, “Emergence of Grounded Compositional Language in Multi-Agent Populations,” in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, S. A. McIlraith and K. Q. Weinberger, Eds. AAAI Press, 2018, pp. 1495–1502. [Online]. Available: <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17007>

- [11] N. Bard, J. N. Foerster, S. Chandar, N. Burch, M. Lanctot, H. F. Song, E. Parisotto, V. Dumoulin, S. Moitra, E. Hughes, I. Dunning, S. Mourad, H. Larochelle, M. G. Bellemare, and M. Bowling, “The Hanabi challenge: A new frontier for AI research,” *Artif. Intell.*, vol. 280, p. 103216, 2020. [Online]. Available: <https://doi.org/10.1016/j.artint.2019.103216>
- [12] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Debiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, R. Józefowicz, S. Gray, C. Olsson, J. Pachocki, M. Petrov, H. P. d. O. Pinto, J. Raiman, T. Salimans, J. Schlatter, J. Schneider, S. Sidor, I. Sutskever, J. Tang, F. Wolski, and S. Zhang, “Dota 2 with Large Scale Deep Reinforcement Learning,” *CoRR*, vol. abs/1912.06680, 2019, arXiv: 1912.06680. [Online]. Available: <http://arxiv.org/abs/1912.06680>
- [13] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou, M. Jaderberg, A. S. Vezhnevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy, T. L. Paine, C. Gulcehre, Z. Wang, T. Pfaff, Y. Wu, R. Ring, D. Yogatama, D. Wünsch, K. McKinney, O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps, and D. Silver, “Grandmaster level in StarCraft II using multi-agent reinforcement learning,” *Nature*, vol. 575, no. 7782, pp. 350–354, Nov. 2019, number: 7782 Publisher: Nature Publishing Group. [Online]. Available: <https://www.nature.com/articles/s41586-019-1724-z>
- [14] T. Rashid, M. Samvelyan, C. Schroeder, G. Farquhar, J. Foerster, and S. Whiteson, “Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 4295–4304.
- [15] C. Yu, A. Velu, E. Vinyals, Y. Wang, A. Bayen, and Y. Wu, “The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games,” *arXiv preprint arXiv:2103.01955*, 2021.
- [16] B. Ellis, S. Moalla, M. Samvelyan, M. Sun, A. Mahajan, J. N. Foerster, and S. Whiteson, “SMACv2: A New Benchmark for Cooperative Multi-Agent Reinforcement Learning,” Jun. 2022. [Online]. Available: <https://openreview.net/forum?id=pcBnes02t3u>

- [17] M. Kim, J. Oh, Y. Lee, J. Kim, S. Kim, S. Chong, and S.-Y. Yun, “The StarCraft Multi-Agent Challenges+ : Learning of Multi-Stage Tasks and Environmental Factors without Precise Reward Functions,” *CoRR*, vol. abs/2207.02007, 2022, arXiv: 2207.02007.
- [18] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [19] A. Michalski, “Informatics Project Proposal StarCraft 2 ”Lite” Environment for Multi-Agent Reinforcement Learning Research,” May 2022.
- [20] F. A. Oliehoek and C. Amato, *A Concise Introduction to Decentralized POMDPs*, 1st ed. Springer Publishing Company, Incorporated, 2016.
- [21] M. Tan, “Multi-agent reinforcement learning: Independent vs. cooperative agents,” in *Proceedings of the tenth international conference on machine learning*, 1993, pp. 330–337.
- [22] C. J. Watkins and P. Dayan, “Q-learning,” *Machine learning*, vol. 8, no. 3, pp. 279–292, 1992, publisher: Springer.
- [23] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. P. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, “Asynchronous Methods for Deep Reinforcement Learning,” in *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, ser. JMLR Workshop and Conference Proceedings, M.-F. Balcan and K. Q. Weinberger, Eds., vol. 48. JMLR.org, 2016, pp. 1928–1937. [Online]. Available: <http://proceedings.mlr.press/v48/mniha16.html>
- [24] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” *CoRR*, vol. abs/1707.06347, 2017, arXiv: 1707.06347. [Online]. Available: <http://arxiv.org/abs/1707.06347>
- [25] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, “Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments,” in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, I. Guyon, U. v. Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, Eds., 2017, pp.

- 6379–6390. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/68a9750337a418a86fe06c1991a1d64c-Abstract.html>
- [26] J. N. Foerster, G. Farquhar, T. Afouras, N. Nardelli, and S. Whiteson, “Counterfactual Multi-Agent Policy Gradients,” in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, S. A. McIlraith and K. Q. Weinberger, Eds. AAAI Press, 2018, pp. 2974–2982. [Online]. Available: <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17193>
- [27] P. Sunehag, G. Lever, A. Gruslys, W. M. Czarnecki, V. F. Zambaldi, M. Jaderberg, M. Lanctot, N. Sonnerat, J. Z. Leibo, K. Tuyls, and T. Graepel, “Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward,” in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, E. André, S. Koenig, M. Dastani, and G. Sukthankar, Eds. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2018, pp. 2085–2087. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3238080>
- [28] J. v. d. Berg, S. J. Guy, M. C. Lin, and D. Manocha, “Reciprocal \emph{n}-Body Collision Avoidance,” in *Robotics Research - The 14th International Symposium, ISRR 2009, August 31 - September 3, 2009, Lucerne, Switzerland*, ser. Springer Tracts in Advanced Robotics, C. Pradalier, R. Siegwart, and G. Hirzinger, Eds., vol. 70. Springer, 2009, pp. 3–19. [Online]. Available: https://doi.org/10.1007/978-3-642-19457-3_1
- [29] Liquipedia, “Starcraft II Liquipedia.” [Online]. Available: https://liquipedia.net/starcraft2/Main_Page
- [30] G. Van Rossum and F. L. Drake Jr, *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam, 1995.
- [31] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del R  o, M. Wiebe,

- P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, “Array programming with NumPy,” *Nature*, vol. 585, no. 7825, pp. 357–362, Sep. 2020. [Online]. Available: <https://www.nature.com/articles/s41586-020-2649-2>
- [32] “Pygame.” [Online]. Available: <https://github.com/pygame/pygame>
- [33] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.
- [34] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [35] Gillies, Sean, Butler, Howard, Pedersen, Brent, Matthias, and Adam Stewart, “Rtree: Spatial indexing for Python.” [Online]. Available: <https://github.com/Toblerity/rtree>
- [36] P. Hart, N. Nilsson, and B. Raphael, “A Formal Basis for the Heuristic Determination of Minimum Cost Paths,” *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968, publisher: Institute of Electrical and Electronics Engineers (IEEE). [Online]. Available: <https://doi.org/10.1109/tssc.1968.300136>
- [37] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, “Cython: The best of both worlds,” *Computing in Science & Engineering*, vol. 13, no. 2, pp. 31–39, 2011, publisher: IEEE.
- [38] S. A. Stüvel, “Python bindings for Optimal Reciprocal Collision Avoidance.” [Online]. Available: <https://github.com/sybrenstuel/Python-RVO2>

Appendix A

SMACLite framework examples

A.1 Example of a valid scenario file

The following is an example custom scenario similar to the built-in scenario `10m_vs_11m`, but using custom units.

```
{
  "name": "10m_vs_11m",
  "custom_unit_path": "smaclite/env/units/smaclite_units",
  "num_allied_units": 10,
  "num_enemy_units": 11,
  "groups": [
    {
      "x": 9,
      "y": 16,
      "faction": "ALLY",
      "units": {
        "example_custom_unit": 10
      }
    },
    {
      "x": 23,
      "y": 16,
      "faction": "ENEMY",
      "units": {
```

```

        "example_custom_unit": 11
    }
}
],
"attack_point": [9, 16],
"terrain_preset": "NARROW",
"num_unit_types": 0,
"ally_has_shields": false,
"enemy_has_shields": false
}

```

A.2 Example of a valid unit file

The following is a custom unit similar to the built-in MARINE unit, but with a much larger (effectively global) scan range

```

{
    "hp": 45,
    "armor": 0,
    "damage": 6,
    "cooldown": 3,
    "speed": 3.15,
    "attack_range": 3,
    "size": 3,
    "attributes": ["LIGHT", "BIOLOGICAL"],
    "minimum_scan_range": 100,
    "valid_targets": ["GROUND", "AIR"]
}

```

A.3 Standard unit types

The following standard unit types are supported: ZERGLING, BANELING, SPINE_CRAWLER, MARINE, MEDIVAC, MARAUDER, ZEALOT, STALKER, COLOSSUS.

A.4 Terrain presets

The following terrain presets are supported: SIMPLE, CHECKERBOARD, NARROW, RAVINE, OCTAGON, CORRIDOR, PENTAGON, ALL_GREEN.

A.5 Standard scenarios

The following standard scenarios are supported: 10m_vs_11m, 27m_vs_30m, 3s5z_vs_3s6z, 2s3z, 3s5z, MMM, MMM2, 2c_vs_64zg, bane_vs_bane, corridor, 2s_vs_1sc, 3s_vs_5z.

A.6 Unit attributes

The following unit attributes are supported: LIGHT, ARMORED, MASSIVE, BIOLOGICAL, MECHANICAL, PSIONIC, STRUCTURE, HEROIC.