Characterizing Human vs Machine Gameplay in StarCraft II

Chad A. Peters (chad.peters@carleton.ca)
Institute of Cognitive Science, Carleton University
Ottawa, ON Canada

Babak Esfandiari (babak@sce.carleton.ca)
Department of Systems and Computer Engineering, Carleton University
Ottawa, ON Canada

Robert L. West (robert.west@carleton.ca)
Institute of Cognitive Science, Carleton University
Ottawa, ON Canada

Abstract

Our research presents a review of the StarCraft II ecosystem, and an analysis of those universal characteristics integral to the replay data generated by thousands of humans and robots in mixed competitions. In this paper we present the obvious and subtle differences between human and machine tournament play, and demonstrate that we can still identify and leverage various aspects of game play to distinguish human from machine.

Keywords: StarCraft; Real-Time Strategy; Behavior Modeling; Cognitive Modeling; Machine Learning; Artificial Intelligence

Introduction

In this paper we describe an analysis of Starcraft game playing replay data to draw attention to differences between human and machine gameplay style, and subtle indicators that may help an observer identify bots that otherwise play very much like a human, and are even capable of defeating expert human players in tournament play.

Our primary motive for replay analysis is less about finding out how to be the best human player, or how to design the best autonomous agent; rather, we are interested to see what sets humans apart from game-playing A.I., and how successfully these aspects of game play can be measured, modeled, and perhaps even used to detect, classify, and replicate these behaviors. We know that A.I. developers describe game playing bots in *terms* of human cognition, and challenge how much of what the bot is and does is actually useful for *understanding* human cognition.

In this paper we attempt to answer these questions through an investigation of game playing behavior produced by humans and machines, both individually, and when playing against each other. The rest of this paper is outlined as follows:

We first review the StarCraft game universe, with an overview of the most recent replay data format and analysis tools, common measurements and their significance, and how this information can be useful to inspect and understand game playing agent behavior in general.

Next, we discusses the human vs human competitive arenas, play styles and metadata that can be used to differentiate players of various skill levels. We find that all human players tend to trend towards a higher and varied rate of effective input as they get better.

The third section deals with machine vs machine tournament play, where A.I. developers can pit their agents against each other in an accelerated tournament environment to evaluate new techniques in intelligent agent design, and specifically for Real-Time Strategy environments such as Starcraft. Our analysis of machine agents demonstrates the tendency of bots to maximize all available action bandwidth provided by tournament servers, and make very little use of the features (or constraints) of a user interface.

In the last section, we discuss the results of recent competitive tournament play between the world's best humans and machines, their apparent similarities and subtle differences in behavior, and some of the controversy involved when A.I. tries to be only as human as necessary. Our results demonstrate that even the most human-like bots are still exploiting non-human abilities in competitions, and possibly disqualifies their use as a model of human cognition.

This paper concludes with a summary of human and robot play styles and indicators, the impact of recent events on the gaming community writ large, and possible future directions for research in this problem domain.

Overview of StarCraft II

The StarCraft II[1] game franchise is a space-opera set in a fictional universe featuring three intelligent racial factions vying for survival and control of limited resources as represented through a series of maps taking place across a variety of terrain. Each of the three playable races (or factions) made available to the player specialize in a unique style of warfare, with corresponding strengths and weaknesses (in Paper-Rock-Scissors fashion) that may appeal to different player preferences.

Real-Time Strategy (RTS) games such as StarCraft require players to successfully balance multiple elements such as resource management, dealing with uncertainty and imperfect information through fog of war, foresight and anticipation, and regularly switching between strategic macro- and tactical

[1] https://starcraft2.com/

micro-management (respectively) to optimize control of various units and groups. RTS games of this nature are known in gaming community and eSports circles for demanding a high level of cognitive performance from players and tend to draw out players with an acumen and appetite for thriving in complex, fast-paced, and high-stakes environments.

Like many RTS games of similar nature, SC2 provides a default set of key mappings that allow a player to quickly input key combinations to accomplish slightly more complex commands. Game command complexity ranges from simple single mouse clicks or keyboard shortcuts, to more complex keystroke combinations for unit group selection or navigating "tech trees" for producing and replacing various units. Command complexity can be measured through a combination of the number of keystrokes required to register a valid selection, and amount of time typically required to complete the input.

The StarCraft community uses different measurements to compare and contrast player performance and aptitude during a match, mostly focusing on the rate of input during different phases of gameplay; example of this include raw input and screen adjustments within a given period of time. The base measurement of Actions Per Minute (APM)[2] can be defined as the lowest level of user input typically associated with the push of a button on the keyboard or mouse. Observation of human replays demonstrates that a high APM, while frequently correlated with a high skill level, does not necessary predict effective gameplay, as many actions are simply repetitive clicking on the same area of the screen. This behavior is perhaps used by some players to maintain a certain micro-management tempo during escalated confrontation. High-frequency actions are not necessarily useful actions, and thus the literature sometimes makes use of Effective APM (EAPM or EPM), to distinguish strings of commands that are both unique and valid from those spammed in repetition.

Sc2gears[3], an online replay analysis site, makes an additional distinction between Micro- and Macro-APM: activities that require resources such as building, training, upgrading, or researching are considered macro-management activities and contribute to overall strategic play, whereas everything else involving individual units or groups for movement and direct engagement with the opponent are considered forms of micro-management.

Screens Per Minute (SPM), another common measurement describing manipulation of the visual playing field, can be defined in a similar way to APM; however, the low level operator in this case is the number of times the player moves the screen in one minute. Moving a screen in StarCraft can take the form of either panning (by using either keyboard or mouse to move the screen in one of four cardinal directions), or by selecting a specific spot on the mini-map.

[2]https://liquipedia.net/starcraft2/APM

[3]https://sites.google.com/site/sc2gears/features/replay-analyzer/apm-types

Figure 1: StarCraft II Game depicting a battle between Terran and Protoss forces

Related Work

Laird and Van Lent(2001) could be credited with one of the earliest initiatives to promote video games as an alternative platform for testing Artificial Intelligence. Different game genres attracted different audiences, from the idle Puzzle Adventure gamer to the hard-core (and somewhat twitchy) First Person Shooter (FPS); as recent history would have it, a combination of balanced and repeatable gameplay and backing from the eSports industry has projected the Real-Time Strategy (RTS) genre into the spotlight, attracting players from all walks of life – for fun, profit, and everything in between.

In (Robertson & Watson, 2014), RTS games like StarCraft have become a de-facto standard for training and testing learning agents, with a growing divide between academic research and the games industry. Researchers are discovering different ways of modeling and understanding spatial-temporal hierarchy problems, however, many papers use different evaluation metrics, making comparison extremely difficult.

Webber et al. (2010) studied players of various skill levels and found those with consistently higher APM usually perform better in RTS games such as StarCraft; their analysis concludes this is due to experts encoding ballistic action sequences. Further, the expert player produces a higher Spatial Variance of Action (distributed attention) yielding a higher probability of yielding required information without causing cognitive overload – professional players know where and what to look for without investing valuable focus time on arbitrary features (Weber et al., 2010).

In (Čertický & Churchill, 2017) we find a review of the current state of these competitions, and the variety of AI bots that compete in them. Growing interest from the gaming industry eventually led to a joint effort between Blizzard Entertainment[4] and Google Deepmind[5] in the creation of a publicly available StarCraft 2 Learning Environment to fast tract

[4]https://www.blizzard.com/

[5]https://deepmind.com/

development of Reinforcement Learning agents, and agents that learn to achieve a level of play that is comparable to a novice player (Vinyals et al., 2017).

Huang et al.(2017) found that some novice and most experts produce excessive APM during the first two minutes of a game; those interviewed associated it with a warm up (not unlike sports games); however, a distinguishing feature between novice and expert is the consistency of APM through the rest of the match – experts rarely decline in APM, whereas novices drop off during periods of intense or confusing states. Further, they also found that experts are more likely to use unit groups for production buildings (vs mobile units), rebind unit groups on-the-fly, and retain consistent habits regardless of game state.

Penney et al.(2018) explored the focus of attention using Information Foraging Theory, and found that while all players have to actively select what to focus on, experts have the highest "return of investment" for their efforts. At a macro-cognitive level, their participants favored *What* information over the *Why* as reported by previous research, and their *Whats* were nuanced, complex, and sometimes expensive, causing participants to dwell on features longer than necessary. They also found players' decision points fell into four main categories of decision cues: building/producing, fighting, moving, and scouting. They found player time spent processing these points were largely dominated by fighting and building, to the point that signs of fighting were classified as major distractor cues(Penney et al., 2018).

Methodology

Our analysis of player matches uses a combination of state-action pairs and metadata derived from StarCraft replay files. SC2 replays are stored in MoPaQ[6] (MPQ) files, an efficient container created by Blizzard Entertainment to store media and gameplay data. MPQ archives can store an arbitrary number of files to encapsulate game state and associated metadata for later retrieval, replay, and in this case, analysis of sequential state-action pairs.

There are various tools to parse the human- and machine-generated replay packs, discover those features most indicative of human players, and to model those features such that they can be compared against their machine-generated counterparts. We primarily used Blizzard's s2protocol[7], a Blizzard python library for decoding SC2Replay files, and Scelight[8], a replay visualization and report generator that excels at build order and ladder career analysis for competitive players.

Game event metadata, such as time stamp (consisting of game loops, each representing approximately 0.0625s), player ID, and command ID (Table 1). The SC2 command taxonomy consists of a relatively simple parent-child hierarchy with the most prevalent commands at the top. Table 1 gives and example of event IDs parsed from a replay file that can be used to categorize classes of user inputs.

ID#	Command
49	Camera Update
104	Cmd Update Target Point
29	Control Group Update

Table 1: Sample command types seen in replay data

All of these tools can be used to extract and visualize various aspects of game play such as current player league standing, command sequences, input (action) frequency, and game events that are presented to each player at specific times. We used this information to derive and compare the common measurements for each combination of human and machine match up as discussed in the rest of this paper.

Human vs Human

The Blizzard ladder league system ranks and matches Human players according to their evolving Match Making Ranking (MMR) score, which in turn is largely based on the Elo Chess Ranking system created by Arpad Elo and adopted by the US Chess Federation[9] to calculate relative skill level between players in organized competition. Blizzard games use Battle.net, a platform-independent system used to match and rank competitive players, and provide an API to access replay data archives associated with each match. The MMR scoring system used by Blizzard is also used to divide players into leagues for general comparison and occasional tournament organization.

Data Sources

In late 2017, Blizzard and Deepmind embarked on a joint effort to create and release a set of tools that could be used to accelerate research and development of intelligent agents in this domain, including a corpus[10] of anonymized human-generated replays for use by researchers wishing to model human players, and for training and testing associate Reinforcement Learning algorithms. Blizzard then released the replays in two sets; the first set is a stratified sample of 64,396 matches, and is a subset of a more complete historical corpus consisting of 1,160,650 replays. Our analysis of human performance characteristics made use of the smaller of the two sets to as it provided sufficient representation of all skill levels (as represented by player Match Making Ranking and placement league), and still feasible for most researchers to reproduce on standard computational resources.

Upon closer inspection of the replay files, we found that each header, although anonymized by player ID, still contained the player's per-match MMR in the replay header metadata. We extracted the MMR along with the average APM per player for each of the 64k replay files to establish a

[6]https://fileinfo.com/extension/mpq
[7]https://github.com/Blizzard/s2protocol
[8]https://sites.google.com/site/scelight/
[9]http://www.uschess.org/about/about.php
[10]https://github.com/Blizzard/s2client-proto#replay-packs

more accurate correlation between observable characteristics and MMR. We then filtered replays due to missing or corrupted headers, incomplete or unknown matches, and missing League information. We resumed with refined corpus of 45,834 replays, each representing two different players, for a maximum of 91,668 unique *plays*.

Replay Analysis

Our overview begins with features that can be easily aggregated, namely APM, SPM, and MRR. The averages across all replays was 90 APM, and 8.42 SPM (Table2).

Player League	Plays	Plays %	Avg APM	Avg SPM
Grandmaster	2,447	2.67%	193	21.64
Master	9,700	10.58%	155	15.68
Diamond	34,526	37.66%	105	9.50
Platinum	18,395	20.07%	76	6.77
Gold	13,308	14.52%	59	5.49
Silver	10,947	11.94%	47	4.52
Bronze	2,316	2.53%	39	3.75

Table 2: Player APM and SPM distribution by League

In Figure 2 we see a moderately positive correlation ($r = 0.65$) between player APM and MMR extracted from each SC2 replay file metadata. This indicates a relationship between the two, with higher average APM likely being the result of player skill level, rather than the cause of it. Breaking out player APM by League standing (Figure3), we observe an increase in both mean and variance of player APM as player skill rises through beginner to expert levels.

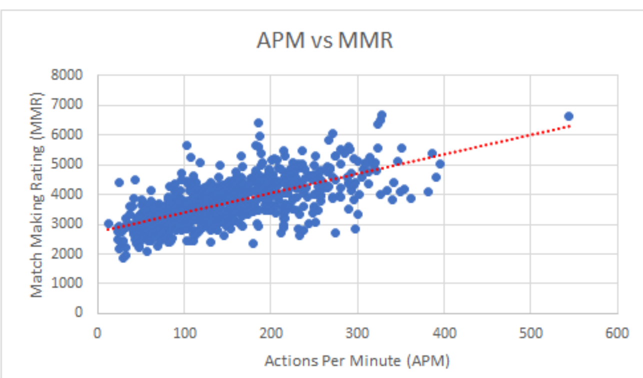

Figure 2: Player APM by MMR; $r = 0.65$

The increase in mean APM is somewhat expected, as players familiar with the game and action sequences will naturally initiate, queue up, and respond to in-game actions on the fly. We also find the increase in APM variance an interesting phenomenon, which could indicate a divergence of macro (strategic) vs micro (tactical) play styles; however, this could also be due to a higher upper bound limited by each player. Each league is normally distributed around the mean with a long tail in the upper-APM range.

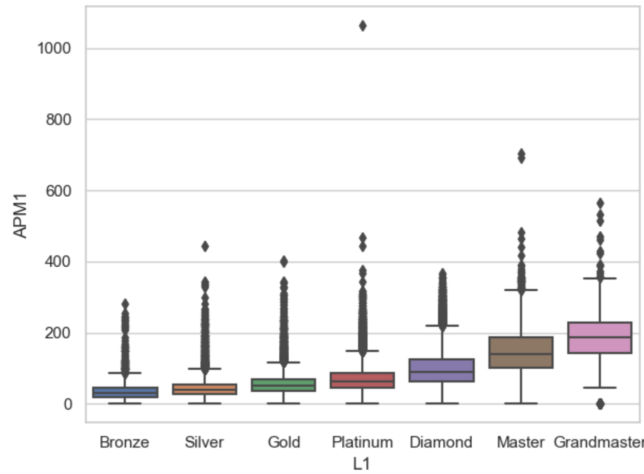

Figure 3: PlayerAPM Quartile Range per League

If we aggregate the average APM and SPM by player league, there is indeed noticeable relationship between league and average APM. We also identified two extreme outliers with one Platinum league player with an average APM of 1064, and two Master league players with 704 and 691 average APM. In both of the later two games we observed cycling between static building control groups multiple times per second. A detailed breakdown of this replay revealed APM spikes as high as 1479 from cycling through control groups; behavior performed at millisecond intervals is only typical of bots observed in the AI arenas, and is considered against the EULA for ranked human vs human Ladder matches.

Machine vs Machine

Bot vs bot gameplay is a popular method to test the efficiency of machine agents in RTS games. We next review are few of the most currently active AI tournament managers that cater to AI research for Starcraft 2.

Data Sources

The SC2 AI Community[11] is one of the most active and longstanding in machine vs machine competitive play, provides ample resources to introduce developers to create agents based on working examples, and also hosts an ongoing Ladder-style matchup service for competitive bot vs bot matches. The StarCraft 2 AI Ladder system continually ranks bots using a similar Elo-based calculation, and can be used as a rough indicator of skill for theoretical league divisions; this ranking system can be used as a rough guide, however we found no evidence of SC2AI bots being divided into leagues as Blizzard does with human players.

[11]https://sc2ai.net/

We first obtained samples of the last few weeks of Season 8 ladder ranking in machine vs machine matches to set a baseline of activity. To do this, we downloaded a minimum of 20 matches per active bot, subject to posted replay availability and game completion, for a total 846 replays. After filtering for crashed games and replay errors, we obtained 798 verified replays. Of these bots, only 10 have made their code available for public inspection, whereas the rest have reserved their right to keep the code private, and only disclose the compiled DLL or Bytecode for offline play. The ability to reserve the ability to choose public or private source code is a feature that allows competitive researchers to participate with a lower risk of losing intellectual property.

Replay Analysis

We found the indicators of interest over all replays differed substantially from our human-only replay results. As seen in Figure 4, the Average APM for all machines was 3465; a roughly 10-fold increase that, while somewhat expected of autonomous agents in this environment, confirms that this machine-only environment is not limited to the same timing constraints as their human counterparts, and will use it when available.

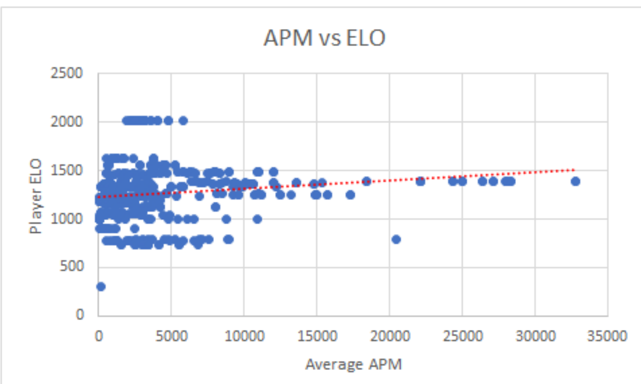

Figure 4: Bots APM by ELO rating; $r = 0.15$

Length range	Plays	Plays %	Avg APM	Avg SPM
0-5 min	47	4.92%	1,830	0.00
5-10 min	426	44.56%	3,926	0.00
10-15 min	346	36.19%	3,726	0.01
15-20 min	71	7.43%	3,306	0.00
20-25 min	24	2.51%	2,208	0.00
25-30 min	4	0.42%	1,641	0.00
35-40 min	8	0.84%	1,431	0.00
45-50 min	30	3.14%	1,025	0.00

Table 3: Bot games by match length: matches have a higher APM by time, negligible use of screen resources.

Also of interest is the almost non-existent use of the screen interface, as reflected by null values for Avg SPM (Table 3). The authors agree that this is likely due to the prolific use of low-level APIs made available to agent developers that allow for access to all permissible game state information regardless of actual location – in essence giving every agent a birds-eye view of the entire match.

Only 5 sessions register Screens Per Minute (SPM), ranging from 0.05 (every 3 seconds) to 0.24 (every 14 seconds). Upon closer inspection there were 4 bots issuing valid Screen Update commands, however, additional investigation is required to determine the causal relationship between observed state and action sequencing without running a code trace.

Correlation coefficient between Elo score and APM is only slightly positive (0.15), indicating the lack of an upper bound on event generation is not a competitive advantage for machines; rather the sophistication of processing (or lack thereof) is more at play. The per-match APM for the majority of bots are between 755 and 3775 with an IQR of 3030, with upper bound outliers going as high as 35,000 APM. This not only captures the majority of games, but the lower bound is well above the peak threshold of professional human players (600), and provides a clear threshold for classification.

Human vs Machine

In the last section, we discuss the results of various tournament play between the world's best humans and machines, derived from a variety of sources. We found that access to quality replays of this nature is somewhat more difficult to obtain, but comparable across game versions using standard metrics as above.

Data Sources

Our first dataset was obtained from the Aritificial Intelligence Starcraft Tournament (AIST) platform[12], that operates tournaments in a hybrid-like fashion. AIST regulations are similar to the well known Student StarCraft AI Tournament (SSCAIT)[13]; however AIST is unique in that they invite high ranking SC:BW players to compete against each seasons' winning bot. We downloaded replays of the final human vs bot matches from the last 3 seasons for a total of 18 replays; all replays were imported without issue.

Our second dataset consisted of a mixture of formal and informal matches played by the well-known AlphaStar[14] agent designed by Google DeepMind. Replay sources consisted of ten (10) public tournament rounds against two WCS champions (TLO and MaNa). The second source consisted of a mixture of replays representing the progression of three learning phases against human players on Battle.net, as discussed by Vinyals et al.(2019).

[12] https://sites.google.com/view/aistarcrafttournament/

[13] https://sscaitournament.com/

[14] https://deepmind.com/research/open-source/alphastar-resources

Replay Analysis

The AIST replays demonstrate a common theme between populations of humans and machines, with a higher level APM vs MMR/Elo of bots. Figure5 contains the aggregate of all tournament rounds between human (H) and machine (M) for the last two seasons, with APM distributions clearly demarcating between the two. AlphaStar, though still quite a bit faster than humans through sustained bursts of high APM, does not make it as obvious through clever use how APM is calculated.

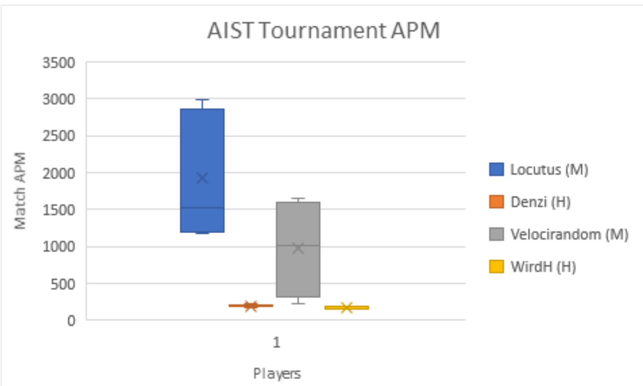

Figure 5: AIST APM for HvM matchup: the two machines have a much higher APM maximum and IQR than their human counterparts.

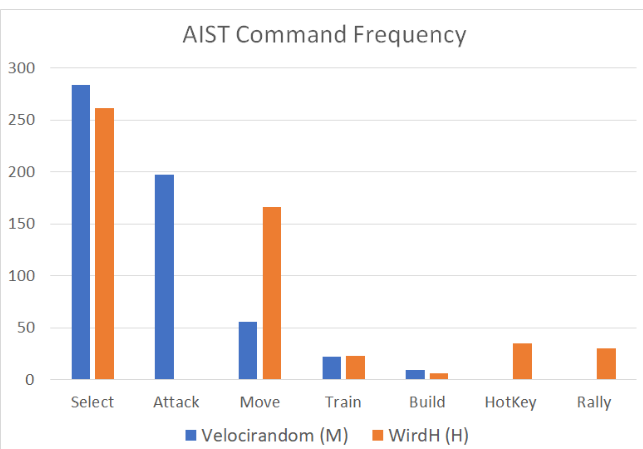

Figure 6: AIST command frequency: humans make greater use of complex commands and unit groupings.

Both replay sources indicate the bots are making direct use of the API, without using a mouse, keyboard, or screen. The ability to see the entire game at once is something of a grey area (but was not against the rules at the time); however, it is still an advantage reserved for agents since human players can only see a small slice of the map at a time. In (Vinyals et al., 2019) AlphaStar tends to focus its attention to one area at a time, which is *sort of* like a human player moving their camera around; however, humans are still at a disadvantage as they still have to move a singular point of focus while using camera controls, at the same time. (Lin et al., 2019) corroborated this with early supervised and mid-tier replays.

With these similarities in mind we took another look at how each of the players interacted with the game itself, and found that all bots, regardless of tournament or game version, do not make use Unit Group Hotkeys (in or out of combat, Figure6). This may appear surprising, however, does make sense as it adds a layer of complexity on top existing data structures, and is only detectable after-the-fact during replay analysis.

Conclusion

In this paper we presented an analysis of StarCraft replay data generated by a variety of sources, to characterize distinguishing features between humans and machines. We provided an overview of the StarCraft tournament ecosystem, and its potentials for both A.I. development and understanding of human cognition.

Our Human vs Human results set the standard for typical performance in raw Actions Per Minute, with is a positive correlation between player APM and skill level in humans, and a few artifacts suggesting some players are trying to use bots or auto-scripting to game the ladder system. This behavior is in stark contrast to the replays evidenced by Machine vs Machine tournaments, where there is a weak relationship between action speed and ranking; agent architectures will take advantage of as much computing resource as possible without consideration for human-likeness or limitations, if they don't have to. Last, we observed a mixed adherence to human-likeness in Human vs Machine match-ups. Despite the shrinking gap between obvious and subtle play styles, in-depth analysis shows us that even the most sophisticated bots will not make use of interface features designed for humans, if they don't absolutely have to.

While some of the results of this research are compelling, there are areas that require additional investigation. We have begun to model game playing events in terms of discrete and continuous stochastic processes through Markov transitions, and would like to continue investigating probable internal representations of behavior using Hidden Markov Models. We hope these and other areas of investigation will shed additional light on what distinguishes a human from a machine while still engaging on common ground.

References

Čertický, M., & Churchill, D. (2017). The Current State of StarCraft AI Competitions and Bots. In *Proceedings of the aiide 2017 workshop on artificial intelligence for strategy games.*

Huang, J., Yan, E., Cheung, G., Nagappan, N., & Zimmermann, T. (2017). Master Maker: Understanding Gaming Skill Through Practice and Habit From Gameplay Behavior. *Topics in Cognitive Science*, *9*(2), 437–466. doi: 10.1111/tops.12251

Laird, J. E., & Van Lent, M. (2001). Human-level AI's killer application interactive computer games. *AI Magazine*, *22*(2), 15–25.

Lin, M., Wang, T., Li, X.-B., Liu, J., Wang, Y., Zhu, Y., & Wang, W.-P. (2019). An Uncertainty-Incorporated Approach to Predict the Winner in StarCraft II Using Neural Processes. *IEEE Access*, *7*, 101609–101619. doi: 10.1109/access.2019.2930581

Penney, S., Dodge, J., Hilderbrand, C., Anderson, A., Simpson, L., & Burnett, M. (2018). Toward foraging for understanding of StarCraft agents: An empirical study. *International Conference on Intelligent User Interfaces, Proceedings IUI*, 225–237. doi: 10.1145/3172944.3172946

Robertson, G., & Watson, I. (2014). A Review of Real-Time Strategy Game AI. *AI Magazine*, *35*(4), 75. Retrieved from `https://aaai.org/ojs/index.php/aimagazine/article/view/2478` doi: 10.1609/aimag.v35i4.2478

Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Georgiev, P., Oh, J., . . . Apps, C. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, *575*, 350–354. doi: 10.1038/s41586-019-1724-z

Vinyals, O., Ewalds, T., Bartunov, S., Georgiev, P., Vezhnevets, A. S., Yeo, M., . . . Tsing, R. (2017). *StarCraft II: A New Challenge for Reinforcement Learning.* Ithaca, New York: arXiv:1708.04782. Retrieved from `http://arxiv.org/abs/1708.04782`

Weber, B. G., Mateas, M., & Jhala, A. (2010). Applying Goal-Driven Autonomy to StarCraft. In *Proceedings of the sixth aaai conference on artificial intelligence and interactive digital entertainment* (pp. 101–106). Retrieved from `http://aaai.org/ocs/index.php/AIIDE/AIIDE10/paper/view/2142`