Title
A Corpus Analysis of Strategy Video Game Play in Starcraft: Brood War

Permalink
https://escholarship.org/uc/item/3f69f8s4

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 33(33)

ISSN
1069-7977

Authors
Lewis, Joshua
Trinh, Patrick
Kirsh, David

Publication Date
2011

Peer reviewed

A Corpus Analysis of Strategy Video Game Play in Starcraft: Brood War

Joshua M. Lewis
josh@cogsci.ucsd.edu
Department of Cognitive Science
University of California, San Diego

Patrick Trinh
ptrinh8@gmail.com
Department of Cognitive Science
University of California, San Diego

David Kirsh
kirsh@cogsci.ucsd.edu
Department of Cognitive Science
University of California, San Diego

Abstract

We present the first ever large scale analysis of actual strategy video game play. Using a corpus of over 2,000 Starcraft: Brood War games from international tournaments, we analyze every player's actions to determine which of their capacities lead to success. We also tie these capacities to their likely cognitive underpinnings, showing that the ability to distribute attention is correlated with winning games. These results have implications for how we might train people to manage critical real world tasks, and for how we approach the project of studying this new medium.

Keywords: video games; starcraft; attention

Imagine if you could play Monopoly without having to roll dice, make change, or calculate the cost of income tax. All you would have to do is take action: buy, build, mortgage. The game would be faster, and you could spend more time thinking about the actions rather than the procedure. Of course this is possible now that Monopoly has been converted into various video game forms, and as such it has joined a long list of strategy video games that provide their players with a rule based environment in which to take action. Free from the implicit constraint that bookkeeping not be excessive, strategy video games enable multitudes of actions in a short time, concurrently with other players, and selected from an almost unlimited selection of possibilities.

In fact, strategy video games can have more in common with real life tasks such as disaster management than with their board-bound forebears. Fires and floods do not wait to take turns in a crisis, and reliable information about their progress is difficult to come by. But unlike these real world tasks, which might happen infrequently or be difficult to collect detailed data on, strategy video games are played constantly around the world and their data are in principle both knowable and recordable.

In this paper we present the first ever large scale analysis of actual strategy video game play. Using a corpus of over 2,000 Starcraft: Brood War games from international tournaments, we analyze every player's actions to determine which of their capacities lead to success. We also tie these capacities to their likely cognitive underpinnings, showing, for example, that the ability to distribute attention is correlated with winning games. These results have implications for how we might train people to manage critical real world tasks, and for how we approach the project of studying this new medium.

In the following sections we relate this novel methodology to previous gaming-related studies.

Relationship to Perceptual Studies

Perceptual load capacity has been studied in video game players by Green and Bavelier, who demonstrated that the processing of distractors in video game players was less affected by high perceptual load than non-gamers (Green & Bavelier, 2006). Similarly, Green and Bavelier (2003) performed experiments testing the visual span and attentional capacity of video game players in comparison to non-video game players. From these studies they were able to show that expert video game players have a higher attentional and visual capacity when compared to a control group. Furthermore, they found that with video game training, non-video game players were able to significantly improve item enumeration in a visual search task under high perceptual load.

Our study is complementary to these studies. While it is very suggestive that video game training can improve perceptual capacities, one might wonder how what players are doing and perceiving leads to this improvement. Studies like ours take a step towards answering that question by analyzing actual game play, rather than relying on before/after tests.

Relationship to Game Studies

The field of game studies consists primarily of game-centered approaches and player-centered approaches (like the Green and Bavelier perceptual studies mentioned above). Game-centered methodologies, including ludology and narratology, involve a qualitative analysis of game content and the players' reactions to the game. Ludological studies treat the game as a set of rules and constraints, while narratological studies treat video games as mechanisms for delivering a narrative (Malliet, 2007). Most quantitative studies of video games involve quantitative analysis of qualitative data: that is, researchers play through video games and code for representations of themes such as violence and gender roles (Brand, Knight, & Majewski, 2003; Smith, Lachlan, & Tamborini, 2003). Notable exceptions include Kirsh and Maglio (1994), who performed a quantitative study of Tetris play by recording keystroke timing and game states, and Douglass (2009), who uses image processing and computer vision techniques to explore video game narrative structure. Douglass relates the interactive/non-interactive screen time ratio to a player's immersion in the game.

The unit of analysis in game studies can include interactions between player and game, not just the game or the player alone. Our study focuses on the interactions between player and game in Starcraft. Additionally, since Starcraft records player/game interactions automatically (see below), we are able to base our analysis on a large corpus of data archived from tournament play. This is in contrast to other models of interface interaction capture, where researchers must bring players into the lab for data collection. Thus our study extends quantitative analysis of game play to an entire population of Starcraft players, and is distinct from studies that focus on the experiences of individual players. Finally, Starcraft is fundamentally adversarial, which allows us to relate actions in the game to actual win conditions rather than more synthetic measures of performance.

A Partial Taxonomy of Games

In order to concretely relate Starcraft to other games and to other challenging activities that are not traditionally construed as games, we will develop a partial taxonomy of games using four fundamental game characteristics. The taxonomy is partial because it is not intended to differentiate any two given games, but rather to illustrate differences between classes of games. The four characteristics are stochasticity, incomplete information, unlimited opportunity and asymmetry. We will define each characteristic and then discuss Starcraft's place in the taxonomy.

Table 1: A short list of games and tasks and their taxonomic characteristics.

Games & Tasks	Stochastic	Incomplete Information	Unlimited Opportunity	Asymmetry
Chess				
Backgammon	✓			
H.H. Hippos	✓		✓	
Mastermind	✓	✓		
Poker	✓	✓		✓
Tennis	✓		✓	✓
Starcraft	✓	✓	✓	✓
Disaster Management	✓	✓	✓	✓
Air Traffic Control	✓	✓	✓	✓
Military Command	✓	✓	✓	✓

Stochastic games might restrict player actions based on the outcome of random events (e.g., one's available moves in Backgammon depend on a die roll), or might similarly modify the game state (e.g., in Risk the outcome of deciding to initiate a battle depends on die rolls). Stochastic games might also modify the game state spontaneously (e.g., winds pushing a ball in Tennis). Of course the behavior of die rolls and wind are the result of deterministic processes, but functionally, for the player, these events are random.

In a game of incomplete information, the game state is not fully available to each player. In poker, for example, one does not know the identity of one's opponents' cards. In Battleship, one knows the identity of the opponent's pieces, but not the position. By contrast, in Chess both players know the complete state of the game at all times. Much like the Markov assumption in machine learning, a Chess player can look at a board and know everything relevant about the state of the game.

Games with unlimited opportunity allow each player to constantly modify the game state through action (limited only by each player's capabilities). In games like Checkers and Chess, one can only modify the state of the game with a single action when it is one's turn. Thus the opportunity to act in those games is limited, both in timing and number. In Hungry Hungry Hippos there are no such constraints—one can constantly attempt to collect marbles with one's hippo, limited only by one's hippo-levering capabilities. Unlimited opportunity is compatible with turn-taking. In Tennis a player must take turns hitting the ball, but he or she can constantly be moving and changing the state of the game. In Chess, though one might be constantly thinking and planning regardless of turn, those thoughts and plans do not change the game state and opportunity is still limited.

Asymmetric games provide players with distinct materials, action repertoires, or territory. Basketball is asymmetric because the players' bodies, equipment and action repertoires are different (though the territory is symmetric). Checkers and Hungry Hungry Hippos are mostly symmetric (barring color, the slight who-goes-first asymmetry in turn taking games, and the stochastic distribution of hippo marbles).

One notable meta-characteristic of a game is whether humans are able to perform optimally when playing it. Given an hour of contemplation and practice, most players can master Tic-tac-toe and never receive worse than a draw. Humans cannot perform optimally in most games, including Chess, Go, Tennis, and Starcraft. Thus while it may be tempting to interpret the above characteristics in terms of what makes a game more difficult (and therefore more interesting to study), that distinction is likely to be meaningless because humans are playing against humans and they are far from optimal play anyways.

We propose, separate from any arguments about difficulty, that games are interesting to study when they share structure with other tasks that are important and that we would like to perform better at. Chess is interesting because of its relationship to formal reasoning and planning (and was thus a darling of GOFAI). Starcraft is interesting because it shares structure with disaster management, air traffic control and military command (see Table 1).

Two concerns immediately arise. First, the taxonomy glosses over important distinctions between, say, air traffic control and Starcraft. There is no explicit adversary in air traffic control, and the materials are planes and pilots, which have significant autonomy. Of course there will be significant distinctions between real world tasks and games[1], but we believe Starcraft has enough structural similarity, and certainly more than, say, Chess, to be worth studying. The second con-

[1]The field of serious games tries to minimize those distinctions so that solving in-game problems solves real problems concurrently.

cern, then, is "Why Starcraft?" What makes Starcraft more appropriate an object of study than any other of the many, many real-time strategy games that share its structure? In the following section we describe Starcraft and show why it is not just *a* game to study in this genre, but *unequivocally the best* game to study.

Starcraft: Brood War

Game Basics

Starcraft: Brood War is a computer game that involves 2 players[2], who manage an armed conflict from an angled top-down perspective. The players use units, similar to Chess pieces, that have different strengths, weaknesses, and capabilities in order to defeat opponent units and destroy opponent structures. Units battle on a two dimensional map, which is completely obscured from the player until revealed by his or her units. Additionally, the player's primary view of the map only spans a small movable portion of the entire playing field (see Figure 1 for a sample screenshot from the game). These design characteristics make Starcraft a game of incomplete information where efficiently distributing one's attention is paramount.

Starcraft players act simultaneously and continuously to accomplish their goals. Starcraft is a game of unlimited opportunity: there is no explicit limit on the number or frequency of actions players can execute to change the state of the game. The units themselves can follow commands with a degree of autonomy. They can navigate obstacles on the way to a target location supplied by the player, or respond when attacked. These behaviors are simple and sometimes unpredictable, so players must manage units constantly in order to ensure desired behavior. This behavior, combined with a random miss chance on some attacks, makes Starcraft functionally stochastic.

In the same way that acquiring money in Monopoly allows one to buy properties and acquire yet more money, economy management is crucial to success in Starcraft. Two resources, minerals and gas, are distributed throughout the map to be harvested by players' worker units. In order to efficiently gather resources from these areas players must build new buildings close to these resources and protect their worker units while preventing their opponent from doing likewise. Players use the resources they collect to build offensive units and buildings that increase their chance of victory. Thus, economy management in Starcraft revolves around securing and harvesting resources and outproducing one's opponent—a difficult task when one is concurrently engaged in multiple skirmishes around the map.

Starcraft's asymmetry is integral to its appeal. Players play as one of three factions, Zerg, Terran, or Protoss, with each faction having a completely distinct set of units and buildings. Games progress very differently depending on the match up, the combination of factions, present in the game. There are

[2]The game supports up to eight players, but almost all professional competition is one on one.

Figure 1: Screenshots from the strategy video game Starcraft: Brood War. Top, the aftermath of a skirmish between Protoss and Zerg units. Bottom, Protoss workers mine minerals and gas in the Protoss player's main base.

six possible combinations, the mirror match ups: Zerg vs. Zerg (ZvZ), Protoss vs. Protoss (PvP), Terran vs. Terran (TvT), and the mixed match ups: Zerg vs. Terran (ZvT), Zerg vs. Protoss (ZvP), Terran vs. Protoss (TvP). A complete discussion of the strategic differences between the match ups is outside the scope of this paper, but suffice it to say that they are substantial enough to warrant separate analysis.

The design characteristics of Starcraft place it in a unique branch of our game taxonomy, distinct from games like Chess and Backgammon, but similar to real world tasks like disaster management (see Table 1).

History, Cultural Impact, and Study

Starcraft was created by Blizzard Entertainment and released in 1998. Its expansion, Starcraft: Brood War, was also released in 1998, and it is the version that is played competitively. Though the game has been around for quite some time it is still quite popular, particularly in South Korea, and still the subject of strategic analysis, from a class developing theory for proficient Starcraft play to an AI competition set in the game environment.

South Korea is host to the world's most developed professional gaming league, which is devoted entirely to Starcraft. Eleven teams each with approximately 15 players compete in televised matches across three 24/7 networks devoted exclusively to esports. While the best players are national celebrities and can earn well north of $100,000 a year in salary, winnings and endorsements, they practice 10 to 13 hours a day, live in team dormitories, and are discouraged from socializing outside the team. This extreme lifestyle has created a population of professional gamers that can execute over 400 actions per minute (APM) in the game. By contrast, a highly accomplished amateur in the United States would likely top out in the mid 200s (beginners start at well under 100 APM). In the context of Starcraft, international play is defined as play anywhere outside of South Korea.

UC Berkeley recently ran a student-led class that modeled advanced aspects of the game. Topics included calculating the effectiveness of unit spatial distribution by calculating the rate at which individual units give and receive damage, and classifying strategies used by professional Korean players through game theoretic analysis (Crecente, 2009).

The Expressive Intelligence Studio at UC Santa Cruz hosted a Starcraft AI tournament in which colleges submitted AI bots to play Starcraft against each other (Expressive Intelligence Studio, 2010). UC Berkeleys winning entry used the computers ability to perform an unlimited number of actions per minute to create a difficult to defeat patterns of unit motion (Huang, 2011).

Replays

Crucial to our analysis is the replay file, a Starcraft feature that allows users to re-watch games after they have concluded. These replay files are records of the actions that both players took and the time that each one occurred. The files are stored in a proprietary binary format, but developers have created software that can decode them in order to catch cheaters and perform basic analyses of individual play. iCCup, the largest organizer of international Starcraft tournaments, keeps a repository of replay files from recent tournaments. The repository allows any player to download and view replays from high-level players around the world, and it allows us to gather a large amount of data easily. Unfortunately some details that would be useful for analysis are not contained in the replay files, such as the amount of resources and units controlled by each player, the positions of units and buildings, and where the players are looking on the map. We must base our analysis solely on actions taken by the players.

No strategy video game has the level of professional and amateur play, availability of data and tools, and depth of understanding through both analysis and experience, that Starcraft has. For these reasons, it is unequivocally the best strategy video game to study.

Methods

We collected 2,302 replay files from Starcraft games played in international tournaments (i.e., outside South Korea) between August 2010 and January 2011 from iCCup. After excluding games where the players did not both have over 75 APM[3], or where the winner could not be determined, our final corpus consisted of 2,015 games.

Building upon András Belicza's Java package for analyzing Starcraft replays to detect cheaters (Belicza, 2011), we wrote analysis software in Java to extract summary statistics of the actions taken in each game. For each game and each player we record the following:

1. Actions per minute (APM), calculated as the total number of actions over game time in minutes.
2. Spatial variance of action (SVA), the 2D spatial variance of all actions with location, such as placing a building or moving units.
3. Macro action count, the total number of macro-related actions. Macro actions are those that help build the economy and production of a player.
4. Micro action count, the total number of micro-related actions. Micro actions are those that manage units during battle, scouting or positioning.
5. Win state, 1 or 0 depending on whether the player won.

We also record the faction used by each player, and the total game time. While it would be nice to record where players were looking on the map and how much material they had at each point in the game, these data are not present in the replay files as mentioned above.

Because each faction plays differently, and each match up is distinct[4], we separate our results out by win state, faction and match up. We perform two-way, one-tailed positive t-tests only between winners and losers within a particular faction/match up combination for both APM and SVA, where the variances are not assumed to be equal. While the same combination of players in a match sometimes appear multiple times in the corpus we feel the games themselves are distinct enough[5] to warrant treating them as independent samples.

Results

Because our games are collected from actual tournament play, the number of samples for each win state/faction/match up combination are different. Table 2 shows the sample size for each combination and Table 3 shows the mean game length in seconds.

[3]We exclude 75 APM and below in order to filter out low skill players and player slots filled by third party observers of the match.

[4]For example, a high APM in one match up might be low for another match up.

[5]They likely occur on different maps, and players must vary their strategies to avoid becoming predictable.

Table 2: Sample size for each win state/faction/match up.

AvB	ZvZ	TvT	PvP	ZvT	ZvP	TvP
Faction A wins	189	49	348	83	165	233
Faction A loses	189	49	348	239	453	256
Faction B wins				239	453	256
Faction B loses				83	165	233

Table 3: Mean game length in seconds for each win state/faction/match up. Note that as in Table 2 the game length for a faction A loss is the same as the game length for a faction B win, and vice versa, so only one is reported.

AvB	ZvZ	TvT	PvP	ZvT	ZvP	TvP
Faction A wins	551	944	749	806	795	932
Faction B wins				819	856	919

Figure 2 (top) shows the mean APM for each win state/faction/match up. All within faction/match up and across win state comparisons are significant at $\alpha = 0.0056$ (after a Bonferroni correction for $N = 9$ comparisons) except TvT, Zerg in ZvT, and Protoss in ZvP. Notably, though, even for comparisons that are not significant, the trend is in the correct direction. If the match up is AvB, then the blue columns show the results for faction A, and the yellow and red columns show the results for faction B, with the winning column on the left within each pair. The ability to quickly perform actions in the game is clearly correlated with success.

Figure 2 (middle) shows the mean SVA for each win state/faction/match up. All within faction/match up and across win state comparisons are significant at $\alpha = 0.0056$ (after a Bonferroni correction for $N = 9$ comparisons) except TvT, Terran in ZvT, and Zerg in ZvP. Similar to APM, the trend is always in the correct direction. A good spatial distribution of actions (and thus attention, see discussion below) is quite predictive of success in most cases.

Figure 2 (bottom) shows the proportion of macro actions (those relating to the economy and production of a player) to micro actions (those related to unit control) for each win state/faction/match up. ZvZ, which is considered a very micro-intensive match up by players, shows a notably lower proportion of macro actions. There does not seem to be a consistent relationship between macro action proportion and winning, however.

Discussion

We find that subjects that are able to most quickly execute actions tend to win. While a high APM could simply be a proxy for practice and experience, we believe it also ameliorates issues posed by functionally stochastic unit navigation behavior. Since the majority of actions are related to unit control, the additional actions allow players to more precisely determine unit behavior. Another way to look at it is that, due to unlimited opportunity, the players who have the higher APM change the game state the most. Since they are trying to change the game state in their favor these changes are likely good for them, unless they make a mistake.

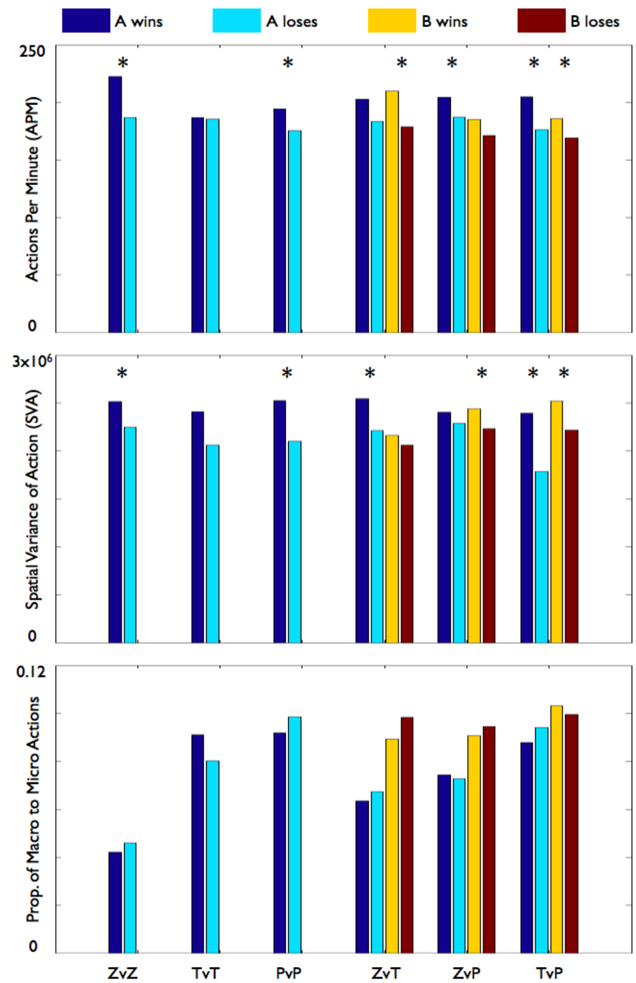

Figure 2: Mean APM (top), SVA (middle) and proportion of macro to micro actions (bottom) for each win state/faction/match up. Within faction/match up and across win state comparisons (dark blue to light blue for faction A and yellow to red or faction B) that are significant at $\alpha = 0.0056$ (after a Bonferroni correction for $N = 9$ comparisons) are marked with an asterisk.

Player SVA is an indirect measure of the distribution of a player's attention. It would be hard for one to take action on a part of the map that one is not at least partially attending to. Given this characterization, we see that the ability to distribute attention and actions around the map (e.g. by scouting in the other players base, or harvesting resources from a different part of the map than one's base) leads to success in

Starcraft. [6] In the tradition of Green and Bavelier, one could imagine testing subjects for multi-tasking capacity before and after Starcraft training. With these data we could then correlate any improvement with actual changes in their attention/action distribution in the games they played for greater insight into the learning process.

Both APM and SVA are useful in Starcraft due to specific components of its design. With a high APM, one can better reduce the stochasticity of unit movement and take advantage of the unlimited opportunity for action Starcraft offers. With a high SVA, a player reduces the amount of uncertainty due to incomplete information by acting on and revealing more areas of the map. In a game that does not have these design characteristics, such as Chess, APM and SVA should have no correlation to win rates. Clearly Chess players do not practice the physical act of moving pieces on the board–there's no need. Even in Speed Chess the emphasis is on quickly developing a strategy and determining the current single best move, not on physically executing a large number of moves in a short time span. If one were managing a disaster, though, high APM should lead to increased performance (e.g. checking in often with workers in the field, issuing instructions, etc.), since the task has similar taxonomic characteristics.

These replay data, which constitute simply a list of actions undertaken by each player, are only a subset of the quantitative data that could be gathered on Starcraft play (though they are the only data suitable for a large scale corpus analysis such as this). With access to players and their machines, one could imagine collecting eye movements, mouse traces, more detailed game state (e.g. changes in material, visual complexity), first person video and audio and hand and body movements. We have developed a simple application to measure visual complexity with first person video, and we would like to continue developing tools to allow a more fine-grained investigation of Starcraft play.

This study develops a novel methodology for analyzing video game play in the wild. Through an analysis of over 2,000 competitive Starcraft games, we show that the ability to change the game state more frequently, and to distribute one's attention and actions around to map leads to success in the game. Since Starcraft shares structure with important endeavors such as disaster management and air traffic control, we believe that a better understanding of what improves Starcraft play can lead to better training for other challenging tasks.

[6] One could argue that a player taking action over a larger area of the map controls a larger area and is thus more likely to win because they are already ahead. While some games of Starcraft are won by building more bases and a larger economy than one's opponent, other common strategies involve strong attacks off of one or two bases while the opponent is attempting to take more of the map. Thus we do not think that the SVA finding is solely the result of map control.

Acknowledgments

Thanks to Darren Lou and Mohamed Mirza for their valuable assistance on this project. Thanks to András Belicza for developing the Java replay analysis package we used to collect the data, and kindly answering our questions.

References

Belicza, A. (2011). *Starcraft Broodwar hacker finder, anti-hack, replay analyzer-organizer and utility tool.* Available from `http://code.google.com/p/bwhf/source/browse/`

Brand, J., Knight, S., & Majewski, J. (2003). The diverse worlds of computer games: A content analysis of spaces, populations, styles and narratives. In *Level up: Digital games research conference.* Utrecht, Netherlands: Digital Games Research Association.

Crecente, B. (2009). *Competitive Starcraft gets UC Berkeley class.* Available from `http://kotaku.com/5141355/competitive-starcraft-gets-uc-berkeley-class`

Douglass, J. (2009). Computer visions of computer games: analysis and visualization of play recordings. In *Workshop on media arts, science, and technology (mast) 2009: The future of interactive media.* UC Santa Barbara.

Expressive Intelligence Studio. (2010). *Starcraft AI competition.* Available from `http://eis.ucsc.edu/StarCraftAICompetition`

Green, C. S., & Bavelier, D. (2003). Action video game modifies visual selective attention. *Nature*, *423*, 534–537.

Green, C. S., & Bavelier, D. (2006). Effect of action video games on the spatial distribution of visuospatial attention. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(6), 1465 - 1478.

Huang, H. (2011). *Skynet meets the swarm: how the Berkeley Overmind won the 2010 Starcraft AI competition.* Available from `http://arstechnica.com/gaming/news/2011/01/skynet-meets-the-swarm-how-the-berkeley-overmind-won-the-2010-starcraft-ai-competition.ars/`

Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, *18*, 513–549.

Malliet, S. (2007). Adapting the principles of ludology to the method of video game content analysis. *Game Studies*, *7*.

Smith, S. L., Lachlan, K., & Tamborini, R. (2003). Popular video games: Quantifying the presentation of violence and its context. *Journal of Broadcasting and Electronic Media*, *47*, 58–76.