

Not All the Bots Are Created Equal: The Ordering Turing Test for the Labeling of Bots in MMORPGs

Stefano De Paoli

Social Media + Society
October-December 2017: 1–13
© The Author(s) 2017
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/2056305117741851
journals.sagepub.com/home/sms


Abstract

This article contributes to the research on bots in Social Media. It takes as its starting point an emerging perspective which proposes that we should abandon the investigation of the Turing Test and the functional aspects of bots in favor of studying the authentic and cooperative relationship between humans and bots. Contrary to this view, this article argues that Turing Tests are one of the ways in which authentic relationships between humans and bots take place. To understand this, this article introduces the concept of Ordering Turing Tests: these are sort of Turing Tests proposed by social actors for purposes of achieving social order when bots produce deviant behavior. An Ordering Turing Test is method for labeling deviance, whereby social actors can use this test to tell apart rule-abiding humans and rule-breaking bots. Using examples from Massively Multiplayer Online Role-Playing Games, this article illustrates how Ordering Turing Tests are proposed and justified by players and service providers. Data for the research comes from scientific literature on Machine Learning proposed for the identification of bots and from game forums and other player produced paratexts from the case study of the game Runescape.

Keywords

Turing Test, bots, deviance, labeling, MMORPGs

Introduction

Recent examples such as the disastrous case of the Microsoft Tay bot and the introduction of chat-bots in Facebook Messenger have brought increased attention to the role of bots in Social Media. According to recent research more than 50% of Internet traffic is generated by bots (Zeifman, 2016). Others have shown that around 15% of twitter accounts are controlled by bots (Varol, Ferrara, Davis, Menczer, & Flammini, 2017). A bot can be defined as “*a program that runs automated tasks over the Internet. Typically intended to perform simple and repetitive tasks, Internet bots are scripts and programs that enables their user to do things quickly and on a scale*” (Gayer, 2016). A bot is a software robot acting in a Virtual Environment (VE) also built on software. A well-designed bot is capable of imitating human behavior and performing actions on behalf of humans. For example, Wikipedia is largely maintained with bots (Geiger, 2009) that support the community by automating repetitive tasks. Bots also have the capacity to shape Social Media online debates (Marres & Moats, 2015). Bots have become a relevant component of Social Media, and it remains an urgent problem to offer sound theoretical and empirical perspectives on their study.

An emerging perspective is traced by Jones (2015), encouraging researchers to focus on the cooperative interactions between humans and bots. Conversely, Jones argues that we should abandon the investigations of the Turing Tests, as these tests focus largely on the functional aspects of bots and have little to offer to the study of social meaning and caring relations between humans and bots. Contrary to this perspective, in this article, I shall argue that the study of Turing Tests is fundamental for revealing social meaning in the interactions between humans and bots. Indeed, not all the human interactions with bots are cooperative as many bots contribute to deviance and crime. Nearly 29% of Internet traffic is generated by “bad bots” (Zeifman, 2016). For instance, Socialbots used for spamming or gathering privacy data have been observed and studied (Boshmaf, Muslukhov,

Abertay University, UK

Corresponding Author:

Stefano De Paoli, Sociology Division, Abertay University, Bell Street, Dundee DD1 1HG, UK.
Email: s.depaoli@abertay.ac.uk



Beznosov, & Ripeanu, 2011). Cheating bots have also been studied in online games (De Paoli & Kerr, 2010). Bots also contribute to building meaning when they are deviant or impact negatively on social order in Social Media. In this article, I will prove that there are sort of Turing Tests proposed by social actors (e.g., end-users, service providers) in which the testing of functional aspects of bots is instrumental to the these actors' capacity of establishing social order. I call these tests Ordering Turing Tests (OTTs), and I will conceptualize them as actors' accounts of how social order can be achieved (Law, 1993). These accounts take the form of what sociologists have called the labeling of deviant behavior (Becker, 1963; Pollner, 1978).

The empirical field of my investigation is Massively Multiplayer Online Role-Playing Games (MMORPGs) as Social Media (O'Donnell and Consalvo, 2015). MMORPGs are widely affected by deviant bots, used in violation of Terms of Services (De Paoli & Kerr, 2010). In MMORPGs, players and service providers craft and justify OTTs which are methods proposed for the labeling of rule-breaking bots and rule-abiding players. I will use two sets of data for researching OTTs. The first is scientific literature on Machine Learning Techniques (MLTs) proposed for the identification of bots in MMORPGs and Virtual Worlds. I consider MLTs as OTTs from the perspective of service providers. The second set of data comes from the case study of the MMORPG Runescape, where I investigated forum discussions on the subject of bots among players and other player produced paratexts (Consalvo, 2007) which discuss bots. The Runescape data shows that players also propose and justify articulated OTTs for labeling bot deviance. The data used in this article are thus the accounts, rationalizations, and justifications that social actors give about their own OTTs.

Love the Bots, Forget the Turing Test

An emerging perspective on the study of bots in Social Media has been proposed by Jones (2015) with a fundamental research question "*How shall we account for social structures that include social machines?*" Jones encourages researchers to accept and "love" the bots for what they are, emphasizing the importance of researching how interactions between human users and bots on Social Media are meaningful in themselves. This appears to be a call to understand bots not as behavioral simulacra of humans or as mere human helpers. Bots and humans entertain authentic cooperative interactions, which have meaning in themselves. An example of this, are the interactions that users entertain with virtual assistants like Siri or Cortana. Jones (2015) argues that "loving the bots" means that "*The [Research] question is no longer whether bots can pass [a Turing Test], but how social interaction with them may be meaningful*" (p. 1). There is a strong emphasis that the investigation of the Turing Test (Turing, 1950) has much to do with functional aspects of bots and thus, says Jones, "*it skirts around the question of the*

social." My perspective on these claims by Jones is that while researching the interactions between humans and bots remains a fundamental research problem, we should however acknowledge that not all these interactions are cooperative and that for understanding cases of bot deviance we should not abandon the question posed by the Turing Test.

The Turing Test is an experiment proposed by Alan Turing (1950) to prove that a machine, like a computer, can exhibit human-like intelligent behavior. The test considers three actors: a human judge, another human (B) and a computer (A). The actors are located in different rooms and communicate via a medium like a teletype. The computer is designed to produce natural language conversations, which are the "*imitation of the behaviour of a man*" (Turing, 1950, p. 435). Based on A and B communications, the goal for the judge is to tell which between them is the human or the computer. If the judge cannot tell the computer apart from the human, then the computer has passed the test thus displaying human behavioral intelligence. This behavioral intelligence, according to some authors (Newell & Simon, 1959), can be explained by specifying computer programs that can produce such behavior and these programs may also offer a way to look at thinking mechanisms with sufficient depth. However, in this article, I am less concerned with "thinking" (whether in humans and/or machines). I rather focus on the interactional aspect of the test and in particular on the judge and her task of telling whether she is interacting with a human or a machine.

My interpretation of Jones's perspective is that the judge in a Turing Test cannot entertain true cooperation with the supposedly intelligent computer. This is because the test aims at making evident the functional limits of the machine in imitating humans. Thus the test does not encourage the judge to care for the machine. This perspective appears aligned with influential studies conducted by Turkle (2007) who showed how kids care for their toy social robots as authentic companions. In a passage of her work, Turkle (2007) gives a brief interpretation of the Turing Test via a parallel with the Voigt-Kampff Test. This parallel reveals many similarities with Jones's propositions.

In the famous Ridley Scott's movie Blade Runner, the main character Rick Deckard is a police officer, whose job is to track down four human-like androids. According to the Law, androids are allowed only in outer-space colonies and not on Earth. However, four of them have escaped colonies, arriving on Earth. Deckard is tasked to find and "remove" these androids. The only way to tell androids apart is to compare them with humans by using the Voight-Kampff Test. The test measures the detection of involuntary emotional responses in the subjects, such as pupil dilation. The core assumption is that human emotional responses are instinctual whereas androids are programmed and the production of a response requires some delay. Toward the end of the story, Deckard falls in love with Rachel, a perfect human-android, and Turkle argues that viewers start doubting whether Deckard himself is

an android. Turkle (2007) suggests that then what really matters is the authentic relation of love and caring between Deckard and Rachel and that

by the time we face the reality of computational devices that are indistinguishable from people, and thus able to pass our own Turing test, we will no longer care about the test. By then, people will love their machines and be more concerned about their machines' happiness than their test scores. (p. 509)

This would seem nearly the same point that Jones makes: tests for telling humans and their robotic companions apart obscure the authentic caring relations which may develop between them.

Apparently, for Turkle the Voigt-Kampff and the Turing Test are the same: mere scores to test robot human-like behavior. While there are similarities between them, it is limiting not to recognize some important differences. The Turing Test is an experiment for testing computer human-like functionalities (Epstein, Roberts, & Beber, 2009). The Voigt-Kampff on the contrary is not a controlled experiment and takes place (at least narratively) in real life. For instance, in one of the opening scenes of the movie a police officer conducting the test is killed by an android. This is not a scenario we expect in a Turing Test. Furthermore, the Voigt-Kampff is a test that some actors (police officers) use for telling humans and androids apart when machine companions apparently violate rules.

To further the understanding of this last point, I reuse an example from the work by Collins (1990) on expert systems. Collins proposed the example of a British spy who needs to pretend to be a native of a foreign country (Soviet Union) from the city of Semipalatinsk. Collins discusses how the spy has learned a great deal of Semipalatinsk from study material and mock-up interrogations. However, the spy has never been in Semipalatinsk. For Collins, the spy is a system that manipulates abstract knowledge about Semipalatinsk, similar to the computer in a Turing Test. Once deployed in the Soviet Union however the spy is captured and is interrogated by a KGB officer. For Collins this setting is similar to that of the Turing Test: the interrogator is tasked with the goal to distinguish between an imitation and a real person from Semipalatinsk. The spy, by manipulating abstract knowledge about Semipalatinsk, can stand an average interrogator. The situation, however, changes dramatically when the KGB brings in an interrogator native of Semipalatinsk. The new interrogator has knowledge which goes beyond the ability of the spy to imitate her origins and can tell the spy apart. What interests Collins (1990) is that "*it will not be possible to construct the equivalent of a socialized being by giving a computer explicit instructions*" (p. 8): a critique to the engineering-functional problem of the Turing Test. However, what interests me is that there is much which is not explicitly said by Collins. It seems implied that the spy is deployed in a foreign country with the intent to disrupt the activities of

the country. Likewise, the spy is interrogated because of apparent violation of the laws of the country. Collins takes the perspective of the spy and the foreign country—how to build a manipulator of abstract knowledge—which is the engineering-functional problem of a Turing Test. For the interrogators, however, the problem is different: proactive testing is needed to tell apart insiders (real citizens) from outsiders (foreign spies). Furthermore, the socialization capacity of the interrogator does not just show the limits of engineering an imitation. Rather, skills and past experiences of the interrogator may play a role in separating insiders from outsiders.

Conceptualizing OTTs

There are sort of Turing Tests—like the Voigt-Kampff or the spy interrogation—for which testing the functional limits of imitations is instrumental for telling apart those imitations who negatively impact social order by not complying with social rules/laws. I call these tests Ordering Turing Tests (OTTs). For conceptualizing OTTs, we need first to define the term social order. I am not referring to a grand narrative of ordering for an entire social system, an eternal social structure. Nor I do refer to a set of rules whose unconditional intersubjective acceptance by social actors delivers order. More modestly, I consider social order as a process enacted by the methodical aspects of everyday life organization, here I would like to focus on a very specific aspect of this: how social actors account, rationalize, and justify for others their decision making and ordering activities (Garfinkel, 1967). We can use Law's (1993) concept of Modes of Ordering to understand the ordering capacities of social actors' accounts and justifications: for Law, actors produce accounts of their actions and these accounts are imputable ordering arrangements, expressions, suggestions, possibilities, or resources for social order. Modes of Ordering are accounts that help actors set boundaries in which elements in a given situation are sorted and labeled. Following this insight, we can formally define OTTs as justificatory accounts and rationalizations of the methods/tests which social actors use for sorting and labeling rule-abiding actors and rule-breaking imitations in real life.

Following this definition, I further conceptualize OTTs as methods for labeling deviance, referring to some elements of the Labeling Theory proposed by Becker (1963). This approach considers how certain actors (e.g., police officers), defined as moral entrepreneurs, are proactive (e.g., because of their formal role) in separating and labeling the insiders—those who are perceived as conforming to the rules—from the outsiders—that are perceived as breaking the rules. For Becker (1963) deviance (or better its designation) is objectified with social construction of labeling where the deviant is not necessarily an individual that has broken a rule but rather "*one to whom that label has successfully been applied*" (p. 9).

Table 1. Types of deviant behavior, adapted from Becker (1963, p. 20).

Types of deviant behavior	Obedient behavior	Rule breaking behavior
Perceived as deviant by moral entrepreneurs	Falsely accused	Pure deviant
Not perceived as deviant by moral entrepreneurs	Conforming	Secret deviant

Table 2. Types of deviant behavior, in the Voigt-Kampff test.

Rule/law: Imitations are not allowed (like androids on Earth)	Obedient behavior	Rule breaking behavior
Perceived as Android by e.g. Deckard	Falsely accused (Human mistaken for an Android)	Pure deviant (Android)
Perceived as Human by e.g. Deckard	Conforming (Human)	Secret deviant (Android mistaken for a Human)

The relevance of Becker's perspective for conceptualizing OTTs remains with its core analytical elements. These elements appear also in the examples of the spy and the Voigt-Kampff test: there is a rule (androids or foreign spies are forbidden), which may be broken by an actor (an individual-android, someone claiming to be from Semipalatinsk), however, whether this violation constitutes deviance (or not) is dependent on whether moral entrepreneurs (e.g., Deckard, the interrogator) apply a label (the individual is an android, the person is a spy). Becker provides a table—which is core to his approach—summarizing the types of deviant behavior where two elements intersect (Table 1): the violation or conformation to a rule done by an individual and the reaction/perception that other social actors have of such behavior. This table shows that there are four possible labeling outcomes: falsely accused, conforming, pure deviant and secret deviant.

The outcomes of a test for telling apart imitations from authentic actors (e.g., Voigt-Kampff, see Table 2) mirror those in Table 1, assuming that the rule states that it is not possible to have "imitations" (e.g., a spy in a country, androids on Earth, a machine behaving like humans) and that moral entrepreneurs are actors like Deckard, the KGB interrogator or even a judge in a Turing Test. The category *falsely accused* is when an individual is labeled by moral entrepreneurs as deviant but in fact this individual has not broken the rule. This would be the case of a human tested with the Voigt-Kampff which is falsely accused to be an android. The category *secret deviant* is when there is a rule violation but moral entrepreneurs fail to label the deviant. This is the case of a machine passing the Turing Test or an android passing as human in a Voigt-Kampff. In Collins, this is when a spy passes as a citizen. The *conforming behavior* corresponds to an individual which the social group label as rule-abiding. In a Voigt-Kampff, this would be the case of a human recognized as a human, whose responses are instinctual and not programmed. Finally, the *pure deviant* is when an individual

breaks a rule and is labeled as deviant. This is the case in which androids are told apart with the Voigt-Kampff or when the socialized interrogator can tell the spy apart. Becker's table makes it clear that rules intersect with the reaction/perception that other social actors (moral entrepreneurs) have about a certain behavior.

An ethnomethodological variation of the Labeling Theory was proposed by Pollner (1978) who observed that, beyond the labels and the reaction of moral entrepreneurs, the methods through which labeling designations are achieved play a fundamental role in the labeling of deviance. Pollner (1978) advanced the proposition that the social group crafts the: "*methodologies through which witches are constituted as detectable entities in the first place*" (p. 271). The labeling methods thus create the conditions of possibilities within which the labeling of the deviant becomes objectified. This objectification requires moral entrepreneurs to produce reflexive accounts of the adequacy of their labeling methods. What is relevant is that this perspective on labeling is aligned with the concept of Modes of Ordering and the definition of OTTs. Indeed, Pollner (1978) emphasizes the ordering relevance of the accounts that actors have for "*establishing and sustaining the response as warrantable*" (p. 280). Focusing the research inquiry on the accounts and justifications of the labeling methods allows social researchers to see how the possibility of error is possible in the labeling designations (Pollner, 1978). In the empirical part of the article, I will show that OTTs, in addition to leading to the analytical four-fold types of deviant behavior (Becker, 1963), are primarily accounts given by moral entrepreneurs who justify the methods they use to reach their labeling designations.

There are qualitative similarities between OTTs conducted in VEs and tests such as the Voigt-Kampff or the Turing Test. Among others, in the Voigt-Kampff, there is the physical co-presence of the moral entrepreneur and the potential "culprit" and the former takes advantage of what he or she sees. In the

Turing Test, the judge and the other actors are separated, and the test requires mediated communication. OTTs conducted in VEs present somehow a mix of both these elements. Immersive VEs are characterized by forms of mediated co-presence (Biocca, 1997), achieved by the medium of an Internet connection and the access of, for example, players to a shared VE populated by avatars. In this mediated co-presence “Users make use of the affordances in the [virtual] environments from which they perceive the structure of the virtual world in ways similar to the manner they construct the physical world” (Biocca, 1997). This perspective recalls the material component of the concept of imagined affordance proposed by Nagy and Neff (2015) where possibilities for action in VEs are based on the interplay between what is perceived by users and the (virtual) materiality (e.g., physical, social) of the environment. Furthermore, the notion of imagined affordance allows to consider that bugs or limits in the design of software (such as androids or bots), can be taken as possibilities for action (i.e., testing). The concept of imagined affordance is relevant for studying how OTTs are conducted in a mediated co-presence. Nagy and Neff see the imagined affordance as a concept that scholars can use to understand the socio-technical materiality of VEs. Differently from them however, I will show that imagined affordances are mobilized by moral entrepreneurs in their accounts and justifications of OTTs.

MMORPGs as Empirical Field

MMORPGs are a genre of computer games, played by a large number of players in persistent Virtual Worlds (Shivan, 2016). Players participate to MMORPGs with avatars, used to interact with the VE, with Non-Player Characters and with other players’ avatars. An avatar starts at level 1, and one of the goals for a player is to increase its level and skills for performing better in the game (Castronova, 2008). Leveling is achieved by accumulating experience points obtained by killing computer controlled monsters. The avatar also improves the performance by accumulating goods such as virtual gold or primary materials (e.g., plants). The gameplay activities of the avatar leveling are, however, often considered repetitive and time consuming. Because of this, bots can be used by cheating player to automate the repetitive aspects of gameplay in what has been defined as automatic-play (De Paoli, 2013). In the most serious cases, bots are used to accumulate virtual goods which can be resold in black markets. MMORPG bots when well programmed also mimic the behavior of human players. Bot owners are at an advantage over fair players since bots allow a much faster avatar leveling, hence creating an unfair competition with players. Game companies regard bots as something that impacts negatively on the service they offer to players. Thus, in most MMORPGs the use of bots is forbidden by legal documents such as Terms of Service (ToS) (De Paoli & Kerr, 2010). Companies therefore need to tell apart bots from human players in order to

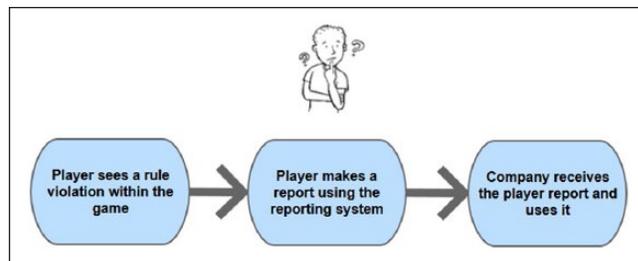


Figure 1. Generic summary of a reporting process in an MMORPG.

apply punishing measures such as bans or deletion of cheating accounts (De Paoli & Kerr, 2012).

MMORPGs constitute an excellent empirical field to study OTTs. I will use two sets of data. The first is computer science literature on MLTs proposed for the identification of bots in MMORPGs and Virtual Worlds. Twenty-three research papers have been collected and analyzed but for reasons of space only eight papers will be cited. Scientific publications can be taken as reasonable representation of the perspective of service providers. What is relevant is that papers publish elements of MLTs which are otherwise not available directly from game companies, which tend to keep secrets on the details of their solutions in order to not offer an advantage to bot makers. Papers are also published by game companies’ researchers (e.g., Lee, Lim, Cho, & Kim, 2016). I will show that MLTs constitute a relevant example of OTTs.

Game companies rely also on the support of players and their reporting (Figure 1) for policing the game from rule violations (Kerr, De Paoli, & Keatinge, 2014). Players are thus encouraged to become moral entrepreneurs (De Paoli & Kerr, 2012). Most games include among reportable rule violations also the use of bots. This leads to having players’ devising their own OTTs before making a report.

To study players’ OTTs, I use data collected from the MMORPG Runescape. Runescape is played over 139 servers each allowing 2,000 players simultaneously.¹ The total number of active players is unknown, but the number of accounts is around 200 million.² The data from Runescape cover a period of time spanning from 2009 until 2014. Most of the data has been collected through the online archives of the game official forums³ and other player produced paratexts, like wiki-pages. Using game forums data is an approach to data collection and research which I used successfully in previous research on MMORPGs and bots (see De Paoli, 2013; De Paoli & Kerr, 2012). In forums and other paratexts players may account for and justify their OTTs. Starting from the end of 2011 the game company has taken an aggressive stance against bots, obtaining relevant successes.^{4,5} Thus, the period investigated is one in which there was fierce competition between players and bots. For this research, 306 forum discussions entirely focused on bots (all those comprising at least 3 pages, at least 30 posts) have been collected and 27

Table 3. Confusion matrix mapped onto Becker's types of deviant behavior.

ML Confusion matrix on rule violation (bots are not allowed in MMORPGs)		
	Obedient behavior (condition negative)	Rule breaking behavior (condition positive)
Perceived as Bot by MLTs (positive prediction)	False positive (FP) (player falsely accused to be a bot)	True positive (TP) (Pure deviant—identified bot)
Not perceived as Bot by MLTs (negative prediction)	True negative (TN) (conforming human player)	False negative (FN) (secret deviant, bot passing for a player)

fully analyzed (all those explicitly discussing tests for telling bots apart). Most of the analyzed discussions include hundreds of pages of posts. In addition, 43 articles produced by the game company and discussing bots have been collected and analyzed. Runescape has been played for four months (2 hours a day) with one character, with the main intent to familiarize with the game terminologies and locations.

Grounded theory (Charmaz, 2014) has been used for data analysis by coding portions of data and subsequently deriving conceptualizations via axial coding. This has supported the development of the conceptual perspective on OTTs described before. Furthermore, to bring both sets of data under the same conceptual frame, the analysis has focused on coding and conceptualizing similarities and differences between MLTs' literature and players' textual data.

MLTs as OTTs

MLTs are described in the literature as solutions to the technical problem of identifying bots in MMORPGs. However, MLTs are entirely predicated on social order concerns. MLTs are OTTs where the functional testing of imitations is instrumental for establishing social order. For example, Kang, Woo, Park, and Kim (2013) justified their novel technique considering that bots “*destroy the game balance by rapidly depleting in-game contents and resources. Honest human gamers may thus feel deprived, lose interest, and eventually leave the game*” (p. 1384). Likewise Mitterhofer, Platzer, Kruegel, and Kirida (2009) justified their technique claiming that bots have a “*severe adverse effect on the day-to-day gaming experience for honest players and impacts their motivation up to the point of making them quit, which has a very real impact on the game company's revenue.*” For reasons like these, the use of bots is forbidden by MMORPGs legal documents (e.g., ToS).

MLTs rely on the concept of User Behavior Analysis: “*the idea that there are differences between human behaviors and programmed bot behaviors*” (Kang et al., 2013, p. 1385). For example, Lee et al. (2016) describe the case of in-game routine activities as follows: “*Our investigations show that game bots frequently repeat certain routine activities that are significantly different from activities of human users.*” The leverage of User Behavior Analysis can be considered an imagined affordance where potential limits in the capabilities of bots (software) to imitate humans constitute the basis

for the labeling of rule-breaking bots against rule-abiding players. Thus, an algorithmic model is trained (supervised learning) on human-play patterns and subsequently the algorithm is tasked to make “predictions” based on this model. Therefore, if a newly observed behavior conforms to the model, the MLT may predict that it belongs to a human. When a newly observed behavior differs from the model, the MLT may predict that it belongs to a bot.

For their functioning, MLTs need to have human behavioral data of game-play available (for training the model) and decisions need to be made on which gameplay aspects the comparison will be made. Where scientists work with or for game companies, they have direct access to the game analytics for a large number of players (e.g., Lee et al., 2016). In other cases, researchers accessed the logs of games (Kang et al., 2013; Mishima, Fukuda, & Esaki, 2013), or recorded the gameplay behavior of volunteer participants (e.g., Gianvecchio, Wu, Xie, & Wang, 2009). From observed user behavioral data, it is possible to compute patterns of in-game human behavior for several gameplay dimensions such as: party-play (Kang et al., 2013), movement on the game maps (Mitterhofer et al., 2009; van Kesteren, Langevoort, & Grootjen, 2009), activity sequences (Lee et al., 2016), mouse/keyboard traffic between client and server (Chen et al., 2008) or user input-actions (Gianvecchio et al., 2009).

The relevant sociological aspect of MLTs is associated with the creation of conditions of possibility for errors in predictions. Supervised learning techniques (not just in MMORPGs) can yield four possible qualitative outcomes: true positive (TP), true negative (TN), false positive (FP) and false negative (FN) (Suthaharan, 2016). These outcomes can be mapped onto what is called a confusion matrix. This matrix is the Machine Learning equivalent of Becker's table on the types of deviant behavior. Turning this insight into MMORPGs (Table 3), we have a rule (e.g., in ToS) which forbids bots in MMORPGs. MLTs are proposed by moral entrepreneurs (scientists, service providers) as methods for separating insiders (rule-abiding players) from the outsiders (rule-breaking bots), based on the imagined affordance of the limits of bots imitation behavior.

The *true positive* case is when presented with a new observation a MLT can successfully tell it belongs to a bot. This is the pure deviant case in Becker. The *true negative* is when upon seeing a new observation the MLT can successfully tell it belongs to a human, the conforming behavior in

Table 4. MLT measures/justifications.

MLT measures (used in justifications of bot detection)	
Accuracy	$\frac{TP+TN}{M}$ Where M is the total number of observations = TP+TP+FN+FP
Precision	$\frac{TP}{TP+FP}$
True positive rate (TPR)	$\frac{TP}{TP+FN}$
True negative rate (TNR)	$\frac{TN}{TN+FP}$
False positive rate (FPR)	1 - TNR
False negative rate (FNR)	1 - TPR

Becker. The *false negative* is one of the cases in which a MLT may make an error. Chen et al. (2008) describe this as “the ratio a bot is mistaken for a human player” (p. 12). This is the Secret Deviant case in Becker (although in MLTs this is a ratio, not an absolute). The *false positive* case is the second case in which a MLT may make an error and is “the ratio a player is mistaken for a bot” (Chen et al., 2008, p. 12). Thus, MLTs as methods proposed to police the game according to set rules (bot are forbidden) contribute in determining the labels, as it is on their user data and data availability, their selected angle of investigation (e.g., action, movement) and so on that we have a labeling, which includes also the possibility for error.

When a supervised learning technique is proposed in any field of application, scientists need to provide evaluations as to the predictive capacity of their models (Suthaharan, 2016). In MMORPGs, these evaluations, while often presented as mere scores, amounts to scientists’ justification of how their OTTs are capable of reducing the possibility of errors in labeling rule-breaking bots and rule-abiding players (see Table 4 for a summary). Most common justifications come in the form of known measures such as accuracy and precision (Suthaharan, 2016). According to Chung et al. (2015) in particular, “Accuracy measures how many bots and humans are correctly identified. Precision measures how many of players detected as bots are really bots” (p. 6).

If a MLT struggles to tell bots apart from humans then it is unable to deliver its intended outcomes. Commercially, it will not be viable, and it will fail to support the policing of a game. Thus, MLTs need to deliver a limited false negative/secret deviant ratio. For example, Chen et al. (2008) using two different schemes for evaluating their technique, show that their “progressive scheme yields a false negative rate of less than 1% and achieves 95% accuracy.” Thus, the false negative rate (1-TPR) is that less than 1 out of 100 bots remains a secret deviant and overall the technique is very accurate (95%) in correctly separating humans and bots.

The false positive rate presents a far more controversial situation because a human could be falsely accused of being a bot. Lee et al. (2016) noted that “False positives should be avoided for the system to be practical. Banning innocent players causes users’ churn, and may raise legal issues and concerns.” Hence, strong justifications of MLTs capacity to

limit false positives/accusations are needed. For instance, Gianvecchio et al. (2009) achieved “true negative rates are 1.0 for all of the humans, so none of the human players in our traces are misclassified as bots.” This means a false positive rate of zero (with FPR = 1-TNR). Chen et al. (2008) showed that their “conservative scheme reduces the false positive rate to zero and achieves 90% accuracy in identifying bots.” This technique achieved zero false positives/accusations (thus 100% precision) but is less accurate (one out of ten bots is a secret deviant). Lee et al. (2016) also claimed that their technique achieves a 100% precision rate, with therefore no false positive/accusation outcome. Thus, while most techniques do seem to be capable of reducing false positive to zero (or near zero), a false positive/accusation remains a possible erroneous outcome requiring strong justifications as to the capacity of MLTs to reduce the possibility of error.

Players’ OTTs in Runescape

Previously, I anticipated that often players are summoned by game companies in a process of reporting bots. The following excerpt from an article written by the Runescape game company explains the reasons for players’ reporting process. The company collects the players’ reports and these contribute to the knowledge base that the company uses for its actions against bots:

Reports for macroing are generally not investigated individually, but instead are added to a “heat map” which Jagex watches closely to find bottled locations and monitor them over time. Jagex uses these reports to help them develop systems-wide solutions to eliminating groups of bots all at once via Botany Bay.⁶

Among players there is a shared understanding that in reporting a bot it is better to avoid false accusations, that is, mistakenly reporting a human player as a bot. This would be unethical and create frictions or unjust punishments. The reporting of bots is based on direct observations of the avatar behavior within the game environment (the mediated co-presence) and on testing. Players recognize that this process presents an intersection between the game rule and the capacity of testers to correctly separate

Table 5. Outcomes of players' OTTs mapped on Becker's table.

Players OTTs' outcomes (e.g., asking something in a chat)	Obedient behavior	Rule breaking behavior
Perceived as Bot by a player-tester	Falsely accused/false positive (Player falsely accused as bot because she does not answer to a chat)	Pure deviant/true positive (Bot does not answer or produces suspicious answers)
Not Perceived as Bot by a player-tester	Conforming/true negative (Player responds to a chat)	Secret deviant/false negative (Bot produces meaningful answers and is taken for a player)

bots and human players. The following excerpt from a wiki page shows this clearly:

*Using bots is not allowed and can be an offence that will ban you. If you know the way to recognize bots, you will be able to report them.*⁷

In their OTTs, players rely on the comparison between what they observe and the supposed behavior that human players should have in the game. Thus, there is a basic similarity with MLTs as also players OTTs may take in account what a tester would expect to be the behavioral differences between humans and bots. In some instances, for players, the identification of a bot can be based on superficial comparisons, for instance looking at the name of the avatar:

*Does it seem computer-generated? Is it just random letters and numbers? Some botting programs change the username, and since it is software, it has little creativity. This step should just raise your suspicion, and you should not take any immediate action until the player is confirmed as a bot.*⁸

The use of random letters to compose an avatar's name is a proof—although not definitive—to tell that there could be a bot controlling the avatar. The underlying logic of this is reflexively justified by the principle that humans would rather use meaningful names for their avatars. However, when bot names are normal, no judgment can be made.

Another simple method is to test whether the suspected bot controlled avatar will respond to a chat:

*Most the time a bot will not respond to your trade request or little "hi!" because no one is at the computer. [. . .] Some modern bots can talk, however. See if they respond suspiciously, such as responding to the same question twice in the exact same way, or responding immediately.*⁹

In this test, the player will ask something using the in-game chat. If there is no answer (or if this is suspicious), the conclusion is that the avatar is controlled by a bot. However, players recognize that this test has limits. Reasons could be that more advanced bots can answer in a meaningful way (i.e., chatbots). Furthermore, many players simply are not interested in answering random chats, as the following excerpt from a forum discussion remarks,

*Just because someone doesn't answer you when you pelt them with questions does not mean they are a bot.*¹⁰

This example shows that the possible outcomes of players' OTTs are the same we have seen in Becker's table and in MLTs confusion matrix (see Table 5). However, too simple tests are accounted by players as producing many errors and not offering sufficient ground for appropriate labeling of deviant (bot) behavior and as a consequence possible of producing false accusations.

Better OTTs for Less False Accusations

Thus, players need better OTTs than those just described for an accurate labeling of bots at the same time avoiding false accusations. Better OTTs are crafted by players exploiting limitations in bots design and taking advantage of the mediated materiality of the VE. In Runescape, prior to the introduction of an update a few years ago, a popular test was the so called Aubury Shop Test¹¹ (AST) for the identification of pure-essence bots. Aubury is a Non-Player Character which can be found in his shop in the Varrock city. Pure-essence is a raw material used to create magic runes. Aubury can teleport avatars to the mine where pure-essence can be gathered. In the AST, the player would take advantage of a limit in the bot design as well as of the game materiality in the forms of buildings and their doors, which can be opened/closed for the purpose of testing.

To access the mine a bot would need to enter the shop in order to be teleported by Aubury himself. The pure-essence bot is thus pre-programmed to reach a specific location (inside the shop building). A player seeing a suspect avatar in the process of entering the shop could then close the door of the shop right before the avatar enters the building. Here is where a bot design limit becomes a testing affordance as players observed that when bots find the door of the shop closed, they try to reach the pre-programmed location by entering a room which is adjacent the shop. However, the adjacent room does not materially allow the bot to access the shop. Thus, when the avatar enters the adjacent room, the player can close the door of this room and trap the avatar inside. If the bot remains stuck in the corner this has proved that the avatar is controlled by a bot (Figure 2). To make sure that there is no false positive, players could also do the following:

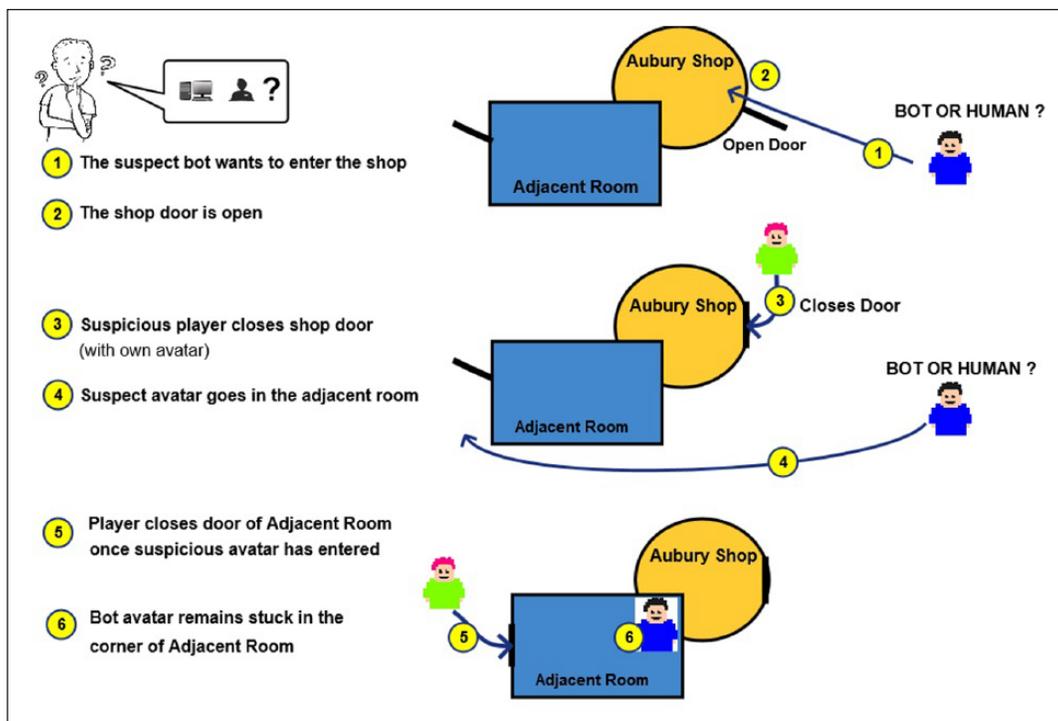


Figure 2. Simplified graphical representation of the Aubury Shop Test.

*To check if it is truly a bot, open the room's door and quickly close it. If the bot runs to the door then runs back to the corner and does so several times again, congratulations you have caught a bot.*¹²

There was also a variation of this test taking advantage of a game item: a cannon or a plant, which could be placed by the tester in front of the Aubury shop entrance. Human players coming back from the mine could walk their avatar past the cannon but players-testers observed that avatars being controlled by bots could not walk past the cannon and would remain stuck inside the shop. OTTs like the AST contribute to minimize false positives/accusations which may be outcomes of superficial testing such as using the chat (Non-talking), as the following excerpt from a post clarifies,

*Non-talking is not a sure sign of a bot, but having them crowd Aubrey's teleport area because of a cannon or a plant, then its kind of obvious.*¹³

Articulated OTTs are both tests used by players for labeling insiders (rule-abiding players) and outsiders (rule-breaking bots), but also are in themselves justifications for reducing the possibility of errors, which are otherwise possible with too simple tests.

Before concluding I would like to present an additional players' OTT to prove some final points. The description of this OTT which can be called Chaos Tunnels Test (CTT) as posted in forums, starts as follows:

Gear Required: Mage robes (Mystic or higher), Runes for teleblock and entangle, Thorny Snail, Food, Melee Weapon/ Offensive spell runes, 1 Click Teleport.

The test requires the player to use specific gears and items which include for instance: food for healing, a weapon and so on. Furthermore, the test requires the player to use a "familiar": a game animal which can be summoned, in this case a Thorny Snail.

The test is done in a location of "the wild," at the entrance of the Chaos Tunnels, which are a series of dungeons. Runescape is organized around safe areas and areas where there is a Player VS Player (PvP) rule and players can attack each other's avatars. This could lead to avatars being killed, consequently losing their gears. The wild is a PvP area. The Aubury Shop instead is located in a safe area. Furthermore, while the AST could be carried out by any type of avatar with no specific level, the CTT requires an avatar with skills level 85 in casting magic spells and 13 in summoning other creatures. The description of the test continues as follows:

Stand one space north of the entrance (so that you are right next to it). Make sure your snail is standing on the far-right space, as this is where most bots enter the tunnels.

Diagram:

.....YT

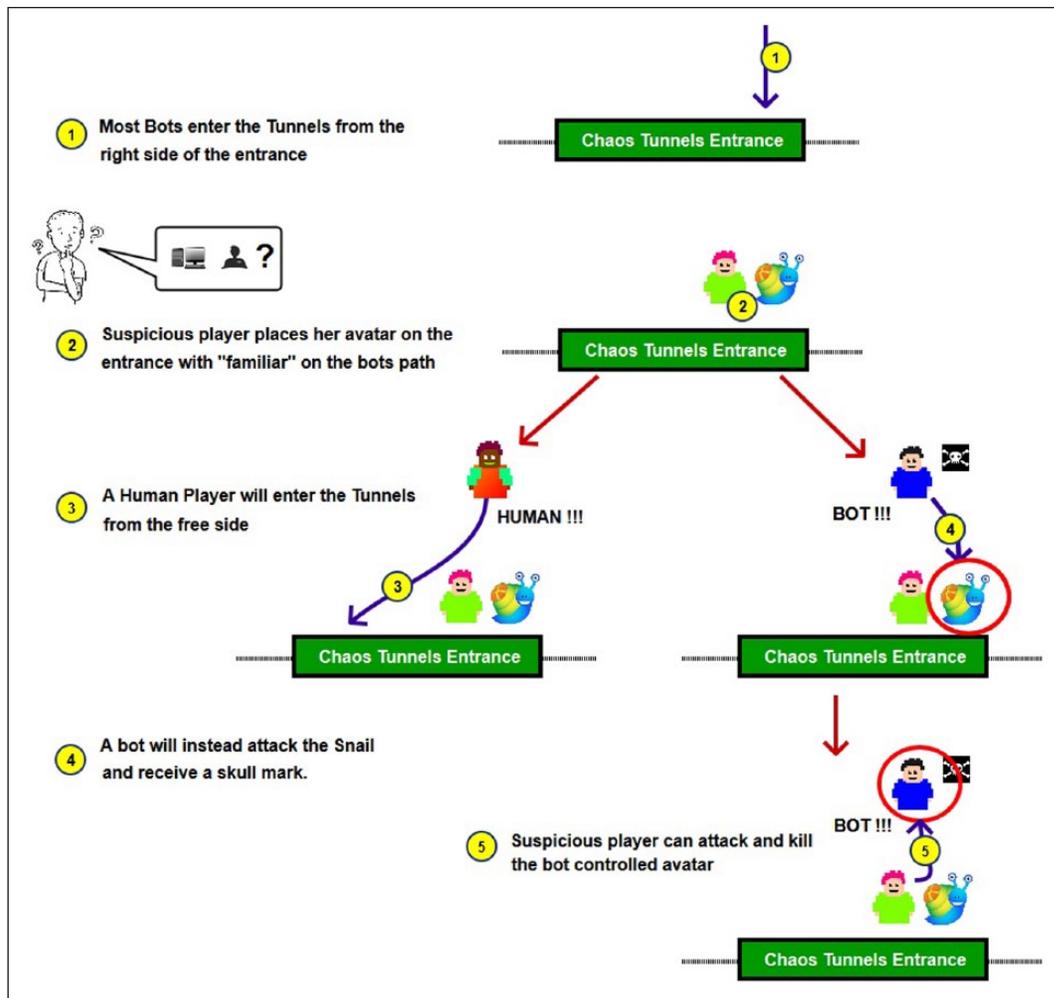


Figure 3. Simplified description of Chaos Tunnels Test.

####

Y= You

T= Thorny Snail

#s = Entrance

Now, just wait for the bot to attack your familiar. If it does not, let it go through and hope to skull it on its next run. If it does skull, entangle/teleblock it and proceed to penetrate its meager defenses with the long stick of your magic/other weapon.

In this OTT, the player needs to place her avatar (Y) at the entrance of the Chaos Tunnels (####) in the manner described in the simplified map, having the summoned familiar (T) on the right side. In the post, the map presents a relevant example of how players mobilize imagined affordances (the materiality of the environment) to offer justifications of their own OTTs. Furthermore, even if this

is not explicitly said, an avatar controlled by a human showing up on the scene and willing to enter the Tunnels can do so by accessing the free side of the entrance. A bot however will find the familiar on its pre-programmed path and will attack it (Figure 3).

The CTT allows the clarification of a couple of points that I made previously. First, like in the case of the police officers killed on duty while conducting the Voigt-Kampff test, players' OTTs are not just experiments to test functionalities in a safe and controlled environment. They are virtual-real life policing processes and thus may also put moral entrepreneurs own avatars at risk. Second, the CTT requires a certain experience (e.g., level 85); this recalls Collins' case of the socialized interrogator. Far from being merely a proof that software cannot imitate humans because it is not socialized, the socialization and experience of the tester is an instrumental factor for the separation of insiders from outsiders and thus for the labeling of bot deviance.

Discussion and Conclusion

This article started by reconsidering an emerging perspective on the study of bots in Social Media (Jones, 2015) which is based on two core claims. First, there is an invitation to Social Media researchers to take seriously the interactions between humans and bots. Second, there is a call to abandon the question of whether bots can pass a Turing Test, because this question obscures the study of social problems in the interactions between humans and bots. While the invitation to take seriously the interactions between (ro)bots and people is crucial for research, in this article I argued for a different position on shifting attention from the question of the Turing Test. I proposed instead that Turing Tests are a fundamental way in which humans and bots interact on Social Media, in particular when bots are deviant and break rules. I introduced the concept of the OTTs to prove that a number of social actors (i.e., users, service providers) propose sort of Turing Tests for the purpose of establishing social order in Social Media.

While there are bots that cooperate with humans (like Wikipedia bots), many other bots produce deviant behavior (like certain Socialbots or MMORPGs bots) and may enter into conflict with humans. Thus, not all bots are created equal. In MMORPGs—the empirical field used in this article—bots produce an unfair competition with rule abiding players and impact on the service delivery of game companies. For these reasons, the use of bots is forbidden by games Terms of Services. When bots produce deviant behavior, social actors resort to using OTTs where the functional testing of bots (the question of the Turing Test) is instrumental for establishing ordering processes in Social Media.

I conceptualized OTTs as ordering accounts (Law, 1993) used for labeling deviance, whereby moral entrepreneurs distinguish between insiders (rule-abiding players) and outsiders (rule-breaking bots). The OTT concept takes elements of two different versions of the labeling theory: the types of deviant behavior from the interactionist version (Becker, 1963) and the focus on the justification of labeling methods from the ethnomethodological version (Pollner, 1978). The main empirical result of this study has been to show that moral entrepreneurs, propose OTTs for separating and labeling bots and humans. Analytically, the elements of OTTs mirror those of the labeling process as described by Becker, as we have a rule (e.g., bots are forbidden in MMORPGs) and we have moral entrepreneurs which judge the adherence of other actors to the rule. This also leads to the types of deviant behavior proposed by Becker: Conforming, Pure Deviant, Falsely Accused, Secret Deviant. OTTs are accounts provided by moral entrepreneurs and the core component of these OTTs is moral entrepreneurs' own justifications and rationalizations as to the reasons why their tests reduce the possibility for errors. Additionally, I have shown that social actors mobilize imagined affordance (Nagy and Neff, 2015) in the justifications of their OTTs.

The concept of OTT can constitute the basis for the study of the interactions between bots and humans when deviance and conflict, more than cooperation, are the core aspects of social meaning and social structures we want to investigate. The presence of deviant bots is not limited to MMORPGs or to Socialbots in platforms such as Facebook. For example, deceptive chatbots pretending to be humans have been observed in dating apps such as Tinder. With this also came the need for users to tell chatbots apart from possible human dates.¹⁴ This is to say that future inquiries will need to expand the study of OTTs to other platforms and other Social Media contexts, with also a focus on building comparative research on bots among different platforms (De Paoli, 2016). Furthermore, since this article has demonstrated that in MMORPGs both algorithmic and players' OTTs, present more similarities than differences, further research is needed to inquiry on whether these similarities are present in other types of Social Media and platforms, for example, Social Network Websites.

To conclude, the main impact of this study could be said to be the refocusing of the debate on the study of the interactions between humans and bots in Social Media. These interactions may very well be cooperative and be based on love and caring as Jones and Turkle argue. However, we need to acknowledge—pretty much like some classical sociologists did (e.g., Simmel, 1904)—that not only cooperation but also conflict is a fundamental component of ordinary social organization. The study of this conflict in relation to bots in Social Media requires appropriate concepts, based on solid empirical research. The concept of OTT has been proposed for this purpose.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Notes

1. <http://en.wikipedia.org/wiki/RuneScape>
2. <http://www.gamesindustry.biz/articles/2013-09-06-runescape-3-boosts-player-numbers-by-300-000>
3. The current forums list is available here: <http://services.runescape.com/m=forum/forums.ws>
4. <http://www.pcgamer.com/runescape-bot-nuking-event-bans-1-5-million-bots-in-one-day/>
5. <http://services.runescape.com/m=news/bot-busting-update-legal-proceedings>
6. http://services.runescape.com/m=rswiki/en/Community_-_Abuse_Reporting_Tips
7. <http://www.wikihow.com/Spot-a-Bot-on-RuneScape>
8. See Note 7.
9. See Note 7.

10. Posted on 29 Oct.2010. Note that all the discussions presented in this article have now been removed from the game forum archives. The author of this article has copies of this data.
11. A partial description of this test and the variation introduced later can be read here: http://darkrunescape.wikia.com/wiki/Door_closer, and here <http://darkrunescape.wikia.com/wiki/Autobuyer>. Videos are also available on YouTube here <https://www.youtube.com/watch?v=vSV9tAbsZPM> and here <https://www.youtube.com/watch?v=mygYJeWIRT4>. The variation using the plant was described in the game forums on 25-Nov-2012. The Chaos Tunnel method was posted on the forum on 01-Sep-2011. A partial description of this test can be found here http://darkrunescape.wikia.com/wiki/Green_dragon_bot.
12. Posted on 16 Apr. 2011.
13. Posted on 12 Nov. 2010.
14. See for an example <http://fusion.net/story/181565/am-i-chatting-with-a-bot/>

References

- Becker, H. (1963). *Outsiders: Studies in the sociology of deviance*. Glencoe, UK: The Free Press.
- Biocca, F. (1997). The Cyborg's dilemma: Progressive embodiment in virtual environments. *Journal of Computer Mediated Communication*, 3(2). Retrieved from <http://www.ascusc.org/jcmc/vol3/issue2/biocca2.html>
- Boshmaf, Y., Muslukhov, I., Beznosov, K., & Ripeanu, M. (2011). The socialbot network: When bots socialize for fame and money. In *Proceedings of the 27th Annual Computer Security Applications Conference* (pp. 93–102). New York, NY: ACM. Retrieved from <https://www.acsac.org/2011/preview/2011-ac-sac-proceedings.pdf>
- Castronova, E. (2008). *Synthetic worlds: The business and culture of online games*. Chicago, IL: University of Chicago press.
- Charmaz, K. (2014). *Constructing grounded theory*. London, England: SAGE.
- Chen, K. T., Jiang, J. W., Huang, P., Chu, H. H., Lei, C. L., & Chen, W. C. (2008). Identifying MMORPG bots: A traffic analysis approach. *EURASIP Journal on Advances in Signal Processing*, 2009(1), 1–22.
- Chung, Y., Park, C. Y., Kim, N. R., Cho, H., Yoon, T., Lee, H., . . . Lee, J. H. (2015). *A behavior analysis-based game bot detection approach considering various play styles*. Retrieved from <https://arxiv.org/abs/1509.02458>
- Collins, H. M. (1990). *Artificial experts: Social knowledge and intelligent machines*. Cambridge, MA: MIT Press.
- Consalvo, M. (2007). *Cheating: Gaining advantage in videogames*. Cambridge, MA: MIT Press.
- De Paoli, S. (2013). Automatic-play and player deskilling in MMORPGs. *Game Studies*, 13(1). Retrieved from http://gamestudies.org/1301/articles/depaoli_automatic_play
- De Paoli, S. (2016). The raise of the robots in virtual worlds: A comparison and a framework for investigating bots in social networks sites and MMOGs. In Y. Sivan (Ed.), *Handbook on 3D3C platforms* (pp. 59–83). New York, NY: Springer.
- De Paoli, S., & Kerr, A. (2010). The assemblage of cheating: How to study cheating as imbroglia in MMORPGs. *The Fibreculture Journal*, 16. Retrieved from <http://sixteen.fibreculturejournal.org/the-assemblage-of-cheating-how-to-study-cheating-as-imbroglia-in-mmorpgs/>
- De Paoli, S., & Kerr, A. (2012). On crimes and punishments in virtual worlds: Bots, the failure of punishment and players as moral entrepreneurs. *Ethics and Information Technology*, 14, 73–87.
- Epstein, R., Roberts, G., & Beber, G. (Eds.). (2009). *Parsing the Turing Test: Philosophical and methodological issues in the quest for the thinking computer*. New York, NY: Springer.
- Garfinkel, H. (1967). *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice Hall.
- Gayer, O. (2016, February 2). Understanding bots and how they hurt your business [Web log post]. Retrieved from <https://www.incapsula.com/blog/understanding-bots-and-your-business.html>
- Geiger, R. S. (2009). The social roles of bots and assisted editing programs. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration* (WikiSym '09). New York, NY: ACM. doi:10.1145/1641309.1641351
- Gianvecchio, S., Wu, Z., Xie, M., & Wang, H. (2009, November). Battle of botcraft: Fighting bots in online games with human observational proofs. In *Proceedings of the 16th ACM Conference on Computer and Communications Security* (pp. 256–268). New York, NY: ACM. Retrieved from <https://dl.acm.org/citation.cfm?id=1653662>
- Jones, S. (2015). How I learned to stop worrying and love the bots. *Social Media+ Society*, 1(1), 1–2. doi:10.1177/2056305115580344
- Kang, A. R., Woo, J., Park, J., & Kim, H. K. (2013). Online game bot detection based on party-play log analysis. *Computers & Mathematics with Applications*, 65, 1384–1395.
- Kerr, A., De Paoli, S., & Keatinge, M. (2014). Surveillant assemblages of governance in massively multiplayer online games: A comparative analysis. *Surveillance & Society*, 12, 320–336.
- Law, J. (1993). *Organising modernity: Social order and social theory*. Oxford, UK: Blackwell.
- Lee, J., Lim, J., Cho, W., & Kim, H. K. (2016). In-game action sequence analysis for game bot detection on the big data analysis platform. In H. Handa, H. Ishibuchi, Y. S. Ong & K-C. Tan, (Eds.), *Proceedings of the 18th Asia Pacific Symposium on Intelligent and Evolutionary Systems* (vol. 2, pp. 403–414). New York, NY: Springer.
- Marres, N., & Moats, D. (2015). Mapping controversies with social media: The case for symmetry. *Social Media+ Society*, 1(2), 1–17. doi:10.1177/2056305115604176
- Mishima, Y., Fukuda, K., & Esaki, H. (2013). An analysis of players and bots behaviors in MMORPG. In L. Barolli, F. Xhafa, M. Takizawa, T., Enokido & H. Hui-Huang (Eds.), *IEEE 27th International Conference on Advanced Information Networking and Applications Workshops* (pp. 870–876). New York, NY: IEEE.
- Mitterhofer, S., Platzer, C., Kruegel, C., & Kirda, E. (2009). Server-side bot detection in massive multiplayer online games. *IEEE Security and Privacy*, 7(3), 29–36.
- Nagy, P., & Neff, G. (2015). Imagined affordance: Reconstructing a keyword for communication theory. *Social Media+ Society*, 1(2), 1–9. doi:10.1177/2056305115603385
- Newell, A., & Simon, H. A. (1959). *The simulation of human thought (RAND Corporation Paper P-1734)*. Retrieved from http://bitsavers.informatik.uni-stuttgart.de/pdf/rand/ipl/P-1734_The_Simulation_Of_Human_Thought_Jun59.pdf
- O'Donnell, C., & Consalvo, M. (2015). Games are social/media (ted)technology too . . . *Social Media+ Society*, 1(1), 1–3. doi:10.1177/2056305115580337

- Pollner, M. (1978). Constitutive and mundane versions of labeling theory. *Human Studies*, 1(1), 269–288.
- Shivan, Y. (2016). *Handbook on 3D3C platforms: Applications and tools for three dimensional systems for community, creation and commerce*. New York, NY: Springer.
- Simmel, G. (1904). The sociology of conflict. *American Journal of Sociology*, 9, 490–525.
- Suthaharan, S. (2016). *Machine learning models and algorithms for big data classification*. Boston, MA: Springer.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.
- Turkle, S. (2007). Authenticity in the age of digital companions. *Interaction Studies*, 8, 501–517.
- van Kesteren, M., Langevoort, J., & Grootjen, F. (2009). A step in the right direction: Botdetection in MMORPGs using movement analysis. In *Proceedings of the 21st Belgian—Dutch Conference on Artificial Intelligence* (pp. 129–136). Retrieved from http://wwwis.win.tue.nl/bnaic2009/papers/bnaic2009_paper_95.pdf
- Varol, O., Ferrara, E., Davis, C. A., Menczer, F., & Flammini, A. (2017). *Online human-bot interactions: Detection, estimation, and characterization*. Retrieved from <https://arxiv.org/abs/1703.03107>
- Zeifman, I. (2016). *Bot traffic report* [Web log post]. Retrieved from <https://www.incapsula.com/blog/bot-traffic-report-2016.html>

Author biography

Stefano De Paoli (PhD, University of Trento) is senior lecturer in Sociology at the University of Abertay in Dundee (UK). His research interests include organizational aspects of cybersecurity, rule-breaking behavior and deviance online and applied user research.