VU Research Portal

Is virtual violence a morally problematic behavior?

Hartmann, T.

published in

Virtual worlds and criminality
2011

DOI (link to publisher)

[10.1007/978-3-642-20823-2](https://doi.org/10.1007/978-3-642-20823-2)

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Hartmann, T. (2011). Is virtual violence a morally problematic behavior? In K. Cornelius, & D. Herrmann (Eds.), *Virtual worlds and criminality* (pp. 31-44). Springer. <https://doi.org/10.1007/978-3-642-20823-2>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Running Head: VIRTUAL VIOLENCE

Is Virtual Violence Morally Problematic Behavior?

Tilo Hartmann
Assistant Professor
Department of Communication Science
VU University Amsterdam
The Netherlands

Abstract

The present chapter pursues the question if virtual violence is morally problematic behavior. Virtual violence is defined as any user behavior intended to do harm to perceived social agents who apparently try to avoid the harm-doing. The chapter reviews existing literature that suggests that users hurt themselves if they get engaged in virtual violence and that users may also do harm other users if they direct their violent acts against their avatars. Recent research is reviewed that further suggests that users may intuitively perceive virtual agents as social beings. They may thus intend to do harm to other social beings (instead of objects) when conducting virtual violence. The chapter also tackles (and denies) the idea that autonomous virtual agents are living entities that may suffer from virtual violence. In light of the reviewed evidence, virtual violence is considered morally problematic behavior.

Is Virtual Violence Morally Problematic Behavior?

Nobody feels guilty about kicking a rock for the simple pleasure of doing so, but doing the same thing to a child is universally forbidden. What's the difference? (Pizarro et al. 2006, p. 82)

Ethics deal with the question if one should or should not engage in a certain behavior. In many cultures violence is considered an ethical issue. Many societies developed norms (and even laws) that tell about justified (or legal) violence and unjustified (or illegal) violence. Norms primarily seek to protect potential victims from harm-doing. However, developing norms about violence does not only protect potential victims, but also potential perpetrators. From a normative perspective that values a socially stable environment, a life of virtue, and the minimization of pain, violence may also harm perpetrators, because it may reinforce or intensify dysfunctional or even pathological personality characteristics (Anderson & Bushman, 2002a; Berkowitz, 1993). Restricting violence does therefore not only help to protect potential victims, but also perpetrators.

Ethical and legal concerns about portrayals of violence in the media are widespread and have been pronounced for decades (Anderson and Bushman 2002b; Geen and Thomas 1986; Slater et al. 2003). Traditionally, concerns focused on detrimental effects of observing violence that is imitated by human-made technology (e.g., in television, music, books). Surprisingly, related discussions and (psychological) theories did not always draw a sharp distinction between the observation of real-world violence and violence in the media. An abundance of studies con-

firmed detrimental effects of observing violence in the media (e.g., Anderson and Bushman 2002b), but the specific implications of observing mediated violence remained somewhat unclear.

The growing popularity of new and interactive media, particularly video games and online virtual environments, revived - and intensified - existing concerns about media violence. In violent interactive environments, people may not only observe violence, but act out violence themselves. Thus, dysfunctional effects on the user may be even greater than the ones observed for non-interactive media like television (Paik and Comstock 1994). Indeed, an abundance of studies conducted within Psychology and Communication Sciences suggest that playing violent video games induces short-term aggression (Anderson 2004). In addition, a growing number of longitudinal studies suggest long-term aggressive effects on the user as well (Anderson et al. 2007; Anderson et al., 2008; Hopf et al. 2008; Möller and Krahé 2009; Slater et al. 2003). In light of these findings, an ethical debate about virtual violence seems to be justified. Accordingly, politicians and lawmakers of key video game countries like the United States, U.K., or Germany started to discuss if access to violent video games needs to be restricted in order to protect users (i.e., perpetrators of virtual violence) from becoming more aggressive.

The present chapter contributes to the existing discussion about the ethical implications of virtual violence (Elton 2000; Luck 2009; McCormick 2001; Powers 2003; Wonderly 2008). Both, users' intention to do harm to others if they engage in virtual violence, and the actually inflicted harm will be considered important aspects when deriving an ethical judgment about virtual violence. The present chapter therefore illuminates if users can be considered intentional aggressors, and if users actually harm themselves and other users (and even other virtual agents) if they engage in virtual violence. In contrast to existing ethical discussions of the topic, the present discussion will be primarily based on social-psychological approaches. Because the existing debate about violence in video games and other virtual environments has not been free of misunderstandings about the subject, a clear definition of virtual violence seems to be a necessary first step.

Virtual Violence

Discussions of virtual violence, i.e., violent behavior in interactive media environments, should be based on a definition of the term. Past definitions of media violence exist, but they do not seem to reflect the specific characteristics of virtual violence.

Some Problems of Past Definitions of Media Violence

Violence (or physical aggression) in the real world is commonly defined as "any behavior following the intention to do harm to others who are motivated to

avoid the harm-doing” (Baron and Richardson 1994, p. 7). The same definition is often used to define *media* violence: „Violent media are those that depict characters intentionally harming other characters who presumably wish to avoid being harmed“ (Anderson et al. 2008, p. 1068). For two reasons, Anderson et al.’s definition may not sufficiently characterize *interactive* virtual violence, however.

First, the definition is about the media, and not about the user. It seems to consider users as *observers* of violence in the media, which is indeed typical for non-interactive media (e.g., television), and may also occur in interactive settings (e.g., video games). However, observation is not a significant characteristic of interactive settings. The significant feature of interactive settings is that users actively manipulate the environment, i.e., they are agents. In violent video games and in other violent interactive settings, users *are* aggressors, rather than merely observers. Users tend to adopt the role of the character they play (Bailenson and Yee 2005; Blascovich and McCall in press; Eastin and Griffiths 2006; Shapiro and McDonald 1992; Yee et al. 2009), instead of merely directing the character like a string-puppet. With more natural interfaces (input-systems and feedback-loops), users’ adoption of the game character’s role may only become stronger. Accordingly, it is a characteristic of virtual violence that it is the user that is intending to do harm (Huff et al. 2003), which should be reflected in a definition.

A second issue that deserves clarification when adapting the definition of Anderson et al. (2008) is linked to the phrase about „characters who presumably wish to avoid being harmed“. One may argue that a character that wishes to avoid being harmed has to be a biologically existing agent, and probably a quite intelligent one, as simple living life forms may not pursue intentional behavior. However, today’s media (or human made technology) are not able to create agents that would match the requirements of biological existence (e.g., Dorin 2004). Accordingly, in a strict objective perspective, media is not capable to give birth to characters that can suffer from harm-doing, or who may wish to avoid being harmed.

Contemporary media technology can only create the *illusion* of living entities. Accordingly, media can only provide the illusion of harm-doing. Eventually, it is up to the user to believe in the illusion of violence (Blascovich and McCall in press). The user must believe that s/he is intending to do harm to a biologically existent agent. Anderson et al. (2008) reflect this notion by speaking of “characters who *presumably* wish avoid being harmed”. Media provides only imitations of reality. Accordingly, media violence is eventually a perceptual phenomenon; media content is only violent in the eyes of a certain user. If users seriously doubt the illusion, however, they may only see pixels on a screen instead of characters, red-colored sprites instead of blood, or well-programmed software scripts, instead of violent actions. If users of violent video games, for example, perceive other things than scared characters and aggressive harm-doing when they move their cross-haired mouse cursor over animated dots of ink that resemble the shape of a human body, it may be misleading to label their action “violent”.

Defining Virtual Violence

A definition of virtual violence should reflect two aspects, then. First, in interactive media settings, violence is not only observed, but users become aggressors as they adopt the role of their game characters. Second, virtual violence requires the perception of others that presumably can be harmed, but apparently try to avoid being harmed. Accordingly, virtual violence may be defined as *any user behavior intended to do harm to perceived social agents who apparently try to avoid the harm-doing*.

Ethical Implications of Virtual Violence

Real-world violence is restricted by law to protect potential victims from being harmed. Virtual violence, in contrast, has become a legal issue primarily because perpetrators, that is, (adolescent) users, may harm themselves (i.e., become aggressive). The discussion in the present chapter will apply an easy but intuitive rationale to ponder the ethical implications of virtual violence. In general, the present chapters starts out from the assumption that violence and aggression are inherently bad (Bufacchi 2004), as they potentially counteract a life of virtue on an individual level, and may dissipate social cohesion and stability of communities on a societal level. One may argue, however, that virtual violence only is morally significant if there is a victim, i.e., a *life form* or a social agent that is *harmed*. Second, focusing on the *outcome* of violent behavior, virtual violence may be considered the more problematic, the more severe the harm-doing to the victim (cf. outcome based ethics). Third, one may also focus on a morally wrong *intention* when judging behavior (cf. duty or intention based ethics). Virtual violence may be considered problematic, because it involves the intention of a user to do harm to other social agents that presumably wish avoid being harmed.

Because the actually amount of inflicted harm is crucial when judging the ethical implications of virtual violence, it seems helpful to differentiate between the (human) user and (virtual) agents as potential victims. Based on such a differentiation, three questions arise. With respect to *users* as victims, one may ask to what extent a user can actually be harmed by virtual violence. Are users hurting themselves if they commit virtual violence as aggressors? (question 1)¹ Virtual violence can be directed against autonomous virtual characters, but also against avatars that belong to other human users. Are other users hurt, if their avatars become victims of virtual violence? (question 2) One may also consider autonomous *virtual agents* as victims of virtual violence. Such a notion seems far-fetched on a first glance, as harm can only be done to living entities. Today's virtual agents can hardly be considered living entities, however². A third question arises that is therefore the most philosophical one: can a (future) media character be harmed? (question 3).

Question 1: Are Users Victims of Virtual Violence?

Virtual violence can be considered a morally relevant behavior, if it has detrimental effects on the (human) user that engages in such behavior.

Users may harm themselves. In the past, an abundance of studies from different scientific disciplines like Social Psychology, Neuropsychology, and Communication Sciences investigated the effects of playing violent video games on aggression (Anderson et al. 2007; Anderson 2004; Bushman and Anderson 2002). Reviews reveal that most studies focused on short-term effects on players. Related studies have shown that playing violent games temporarily increases a player's aggressive thinking, feeling, behavioral intentions, and even behavior (e.g., Anderson 2004). An increasing number of longitudinal studies also examined the effects of repeated violent game play on the personality of players (e.g., Anderson et al. 2007; Hopf et al. 2008; Möller and Krahé 2009; Slater et al. 2003). The existing longitudinal studies suggest that more aggressive people prefer violent video games more than less aggressive people, and that playing violent games reinforces and may even increase aggressive personality characteristics.

In light of the converging evidence that playing violent games increases aggression, it appears that users harm themselves if they engage in virtual violence. The *likelihood* of getting harmed by engaging in virtual violence is comparable to the likelihood of harm resulting from other behavior, such as the risk to get lung cancer from passive smoking at work (Bushman and Anderson 2001). However, the problematic outcome of virtual violence (i.e., becoming aggressive) may be considered less dramatic than the outcome of passive smoking (i.e., getting cancer). Nevertheless, virtual violence appears to be morally problematic behavior, because it seems to affect users in problematic and harmful ways.

Two critical remarks may be made, however. First, problematic outcomes of virtual violence have been primarily examined in the past in the context of playful environments like video games. Little is known so far if the observed aggressive effects apply to other virtual environments as well (McCall et al. 2009). Second, the present chapter defined virtual violence as harm-doing against *perceived social agents*. Past research has shown that users of violent video games may become aggressive. But most studies did not take users' subjective perceptions into account. Therefore, it is not totally clear if the observed effects on aggression are indeed due to virtual violence as it is defined in the present chapter, i.e., users' intentional harm-doing against *perceived social agents* who apparently try to avoid the harm-doing.

In order to ethically judge not only violent video game play, but virtual violence, a closer look on how users' perceive other agents and violent situations depicted in virtual environments seems to be necessary. Exploring how users perceive virtual violence not only promises to illuminate problematic outcomes of virtual violence, but may be also particularly helpful in revealing potentially problematic intentions of users.

Users may follow problematic intentions. Only a few studies exist to date that shed light on how users subjectively experience virtual violence. Neuropsychological research conducted by Mathiak and colleagues (see Mathiak this volume;

Weber et al. 2006; Mathiak and Weber 2004) shows that users' brain activity during violent video game play resembles the brain activity one would expect in similar situations in the real world. In their study players suppressed their emotions if they engaged in violent actions in a video game, i.e., they tried to "cool down". Results further suggest that the activity of brain regions that are important for empathetic feelings toward others was diminished during acts of virtual violence. In sum, the findings of Mathiak and colleagues indicate that users respond in a quite natural way on violence depicted in virtual environments.

Hartmann and colleagues (Vorderer and Hartmann in press; Hartmann et al. in press) suggest that virtual characters are getting increasingly less artificial and may thus develop into "something" a user cannot easily discard as an object, but may rather intuitively consider "some sort" of social and moral entity (Dorin 2004; Elton 2009). Referring to the initial quotation by Pizarro et al. (2006), violence against today's agents that are populating virtual environments may have less in common with kicking a stone, but may be more appropriately understood as a form of interpersonal aggression against another social being. In accordance to social-psychological dual-process models (e.g., Epstein and Pacini 1991; Smith and deCoster 2002), Hartmann et al. (in press) suggest to differentiate users' spontaneous (automatic) perceptions from their more reflective (elaborate) perceptions of violence. Upon conscious reflection, users are likely to be aware of the artificial nature of media violence. Therefore, they probably consider the depicted violence to be "just a game" or to "have nothing to do with the real world" (Klimmt et al. 2006; Ladas 2002).

However, users' *automatic perceptions* of violence may evoke another sensation. In two studies, Hartmann et al. (in press) explored guilt as an indicator of users' automatic perceptions while playing violent games. Guilt can be defined as "the dysphoric feeling associated with the recognition that one has violated a personally relevant moral or social standard" (Kugler and Jones 1992, p. 218). Guilt should indicate if players recognize to violate a moral or social standard if they engage in violent video game play. If users would discard video game violence as completely artificial while playing (as many avid players claim upon conscious reflection), guilt would be an improbable response. Why should anybody feel guilty for "removing a sprite from a computer screen"? Players would not feel to do something wrong, if they were fully aware that they do not hurt somebody, but are actually only immersed in a illusion. In such a case, guilt is an unlikely response. In addition, guilt should also not vary between different levels of justification of the conducted violence, nor should it relate to other determinants known to affect guilt in the real world (e.g., violence against innocent children vs. violence against male soldiers). To the extent, however, players intuitively perceive their violent actions (and the victims) to be meaningful, if not to say: real, they may feel guilty for what they do, especially if the conducted violence is considered unjust and if players are rather empathetic persons.

In the two studies conducted by Hartmann et al. (in press), justification of video game violence and users' trait empathy indeed determined guilt in a structurally

similar way as in real-world scenarios: people felt guiltier if they engaged in unjust violence, especially if they were empathetic persons. Similar to the conclusion drawn by Mathiak and colleagues, the findings suggest that, to a certain degree, players fail to acknowledge the artificial nature of the depicted violence.

A plausible explanation of this result is that users' automatic processes precede their more reflective processes while using a virtual environment (Bargh 1994; Smith and deCoster 2002). Automatic (i.e., intuitive, spontaneous, impulsive) responses are quickly triggered by the media environment, and can only be subsequently regulated upon conscious reflection (Zillmann 2006). Reality-fiction distinctions, which are key in the separation of the media world from the real world, apparently require operations of the reflective system (Smith and deCoster 2002; Zillmann 2006). In contrast, the automatic system seems to be incapable to judge what is mediated and what is not; accordingly, everything that is automatically perceived is subjectively real. Automatic processes seem to quickly form the impression of an apparent reality, even if the sensory input was created by media or other human-made technology (Dorin 2004). The intuitive perception of virtual agents as social agents seems to accompany such an apparent reality, just like emotions that follow on automatic appraisals of the environment (Roseman 2001). Even moral judgments may be intuitively processed (Haidt 2001). Thus, feelings of guilt may have already been triggered before more conscious or reflective elaborations about the displayed media environment set in. Accordingly, automatic social perceptions, emotions, and moral judgments may all already exist, before users are even able to blend in their knowledge that "this is not real" or "this is just a medium".

Editing or regulating perceptions, emotions, and moral judgments once they occurred is, in psychology terms, quite effortful. It may be doubted, if media users are able and actually want to fully regulate or discard their automatic experiences (e.g., Harris 2000). First, a constant flow of incoming sensory information fuels the automatically established apparent reality, which in turn constantly contradicts the more reflective notion that "this is not real" (Meehan et al. 2002). Second, in the heat of a video game play or while being engaged in other interactive environments, users may lack the cognitive resources to fully regulate their automatic impressions. Third, users may only rarely ever be motivated to continuously discard their sensation of an apparent reality. Moviegoers, book readers, video gamers, and users of other virtual environments all rather seem to enjoy believing in the automatically created illusion for a while (Harris 2000). Indeed, research on media entertainment suggests that the enjoyment of most media offerings is much higher, if users are able to believe in their apparent reality (Green et al. 2004), whereas reminders of the artificial nature of the depicted environment can be quite annoying (e.g., a program error, a badly designed character, a totally implausible behavior, etc.).

Taken together, it seems plausible that users of virtual environments may indeed engage in virtual violence as defined in the present chapter, i.e., they may indeed have the (automatic) feeling to do harm to another social agent that seeming-

ly tries to avoid the harm-doing (Mathiak and Weber 2006; McCall et al. 2009; Dorin 2004). To the extent users intentionally pursue their behavior, even though they believe they are harming a social agent instead of an object or “pixels on the screen”, their intention seems to be morally problematic (Powers 2003). Intending to hit another virtual agent does not seem to be the same as the intention to hit, for example, a figure in a chess board game, then. The intention to hit a virtual agent resembles the precarious intention of real-life aggression much more than the harmless intention to hit a chess figure.

Question 2: Can Other Users Be Victims of Virtual Violence?

Many virtual environments feature mediated user-to-user interactions (cf., Konijn et al. 2008). A user directs his or her avatar to interact with another avatar that is navigated by another user. Both users may feel to communicate to each other, if their avatars perform symbolic behavior. User-to-user interactions that are processed through the virtual cloaks of avatars can turn violent as well, if users try to harm the avatar of another user that wishes to avoid being harmed (e.g., Hunter and Lastowka 2003). The most severe form of harm-doing against another user is referred to as “online player killing”, in which a player eventually kills another user’s character inside the game world (Whang and Chang 2004, p. 596). The question arises to what degree users may feel harmed if their avatar becomes a target of virtual violence and eventually dies in the virtual environment? Research that examined this question is growing. Most publications, however, entail theoretical discussions and only include anecdotal evidence of harm-doing, in place of systematic empirical examinations.

Suler and Phillips (1998) provide a detailed review about different deviant behaviors in virtual environments. Based on anecdotal evidence, their report also describes hate and violence avatars that may make use of offensive language. Effects on potential victims are not examined, however. Powers (2003) provides a philosophical discussion of the question if “the harms that some people claim to have suffered in cyberspace [are] real moral wrongs?” (p. 192). Based on speech act theory, he argues that virtual performances are a part of authentic social practices. As such, they are intentional, meaningful, and thus often powerful enough to hurt other users. “For these reasons, it makes sense to speak of moral patients as having suffered real moral wrongs, and accordingly to assign blame to moral agents for having committed these wrongs” (p. 192). However, according to Powers, virtual harm-doing may still not be morally problematic if it adheres to the (explicit or implicit) set of rules of virtual environments. Just like it is not considered unfair if a boxer hits his opponent in the face in an official box fight (e.g., Bredemeier and Shields 1986), virtual violence may be a fair action if it does not violate the rules of the virtual environment.

Huff, Miller and Johnson (2003) discuss a virtual rape that occurred in the text-based multi-user environment LambdaMOO. They come to the conclusion that “perhaps the most important lesson is that virtual actions and interactions have consequences for flesh-and-blood persons and hence, the flesh controllers of vir-

tual action, whether they control directly (as in playing a character) or indirectly (as in designing a virtual world), have responsibilities for their actions” (p. 19).

A virtual rape is an extreme example of virtual violence, but even more common events than virtual rapes may cause harm to other users. In a study about the very popular Massive Multiplayer Online Role-Playing Game (MMORPG) “Lineage”, Whang and Chang (2004) note that a common event like a virtual death can be “a traumatic experience for some players, since dying comes with severe penalties with various consequences such as losing valuable player items or damaging the abilities of players, causing time losses” (p. 596). However, similar to Powers (2003), they concede that there are also legitimate forms of online player killing, if the action adheres to the official rules and norms of the environment. Whang and Chang assume that a “fair killing” of another avatar may also lead to less traumatic user experiences.

Wolfendale (2007) stresses that users’ *attachment* to their avatars may strongly influence how much pain is caused if the avatar is hurt or even killed. Wolfendale’s bases her approach on the view that “avatar attachment is expressive of identity and self-conception and should therefore be accorded the moral significance we give to real-life attachments that play a similar role” (p. 112). Many players feel attached to their characters. As people that seek enjoyment, they are also emotionally engaged in the online world. As a consequence, they may easily feel hurt if their character is harmed. Therefore, especially virtual violence against avatars of highly involved users may be morally problematic: “Attachment to people, possessions, ideals and communities means that we suffer when these are harmed. If we accept such suffering as the normal human condition and as the price we pay for the joy that attachment can bring us, then there is no reason not to accord avatar attachment the same moral standing common experience for many participants who are emotionally engaged in the online world” (p. 118, see Dorin 2004, for a similar notion).

In sum, it appears that other users *can* be harmed if their avatar becomes a victim of virtual violence, although systematic empirical evidence is lacking. The reviewed anecdotal evidence suggests that users feel hurt if their avatar is harmed, the more they feel attached to their character, the more the harm-doing violated existing environment-based norms and related expectations, and the greater the loss of time and monetary resources that have been invested in the virtual character. Because it is wrong to hurt another user, it also seems to be wrong to hurt or even kill another avatar, if the user of an avatar felt emotionally attached to the virtual character, if the harm-doing violated game- or environment-specific rules, and if the killing implies the loss of a reasonable investment of time and money. Accordingly, virtual violence that is directed against other users’ avatars may be morally problematic behavior because it may hurt other users as well.

Question 3: Can a Media Character Be Harmed?

Today’s media characters can probably not be regarded living entities (e.g., Boden 2000; Dorin 2004; Floridi and Sanders 2004; Ray 1992). Therefore, they cannot be harmed. But it seems a worthwhile philosophical question if human-

made technology, specifically computer technology, may be capable to generate life forms, i.e., living virtual creatures, one day (a related discussion may also shed some light on the life status of today's virtual characters). Eventually, today's media technology may only be considered an intermediate step on the way to more advance technologies. A brief look into the discussion about computer-generated artificial life therefore helps to clarify if virtual creatures may qualify for potential victims of virtual violence one day.

Since the rise of computer systems, a lively debate evolved about the possibility of *artificial* life (so-called "Alife", e.g., Searle 1980) and particularly virtual artificial life ("Soft Alife"; e.g., Ray 1992).³ Life has been characterized in many ways, but *self-organization*, including metabolism, emerged as a core criterion of life in many related discussions (Boden 2000). Skeptics like Boden (2000) argue that virtual artificial life is impossible, because electronic systems (e.g., software programs or simulations) lack a physical body that would be able to metabolize in a strict biological way: "Metabolism [...] is the use of energy-budgeting for autonomous bodily construction and self-maintenance, and no actual body construction goes on in simulations of biochemistry" (Boden 2000, p.122). Other researchers, however, stress the evolutionary underpinnings of life and argue that software programs may evolve – and thus may indeed come to life – within a digital environment (Ray 1992; Spafford 1989). Computer viruses, for example, have been discussed (but also rejected) as a form of artificial virtual life, as they resemble a self-replicating organism (Spafford 1989). Similar discussions evolved around whether or not a virtual agent (e.g., software) may develop into a moral agent one day (e.g., Anderson and Anderson 2007; Floridi and Sanders 2004), and if software systems can develop consciousness (e.g., McDermott 2007) or are capable to think (Searle 1990).

A full review of these and related discussions goes certainly beyond the scope of this chapter. It appears, however, that skeptics that doubt the possibility of artificial virtual life (Soft ALife) outnumber the researchers who propose that virtual life may be possible. If the skeptics are right, ethical considerations of virtual violence do not need to consider any actual harm that is inflicted upon a virtual agent, simply, because virtual agents are not really alive: not today, but probably also not in the future.

Conclusion

The present chapter aimed to contribute to the discussion about ethical implications of virtual violence. Related social-psychological literature (and literature from neighbored disciplines) was reviewed in order to illuminate virtual violence, and particularly the harm-doing that results from such behavior. The reviewed evidence suggests that users risk to hurt themselves if they enact the role of an aggressive character, because they likely become more aggressive (and thus support socially dysfunctional states and personality characteristics). From this perspective, virtual violence can be considered ethically problematic behavior.

The reviewed literature suggests that users may also harm other users if they hurt their avatars. The existing empirical evidence, however, is much scarcer

compared to the abundance of studies that proved aggressive effects. The reviewed literature suggest that other users may only be harmed if they felt emotionally attached to their avatars and thus suffer, if their avatars are hurt or even killed. The killing of an avatar may be especially distressful for users, if it implies the loss of considerable time and money. The killing of an avatar, even if it induces distress in the other user, may be still a just behavior, however, if the violence adheres to the (official or unwritten) rules of the virtual environment (Bartle 2009). Such an argument would be analogous to the justification of violence in sports (e.g., boxing), which certainly also implies distress and a loss of invested time and money for those athletes that are hurt (Bredemeier 1985). Still, the violent action may be considered appropriate, as all athletes willingly entered the violent situation and thus were able to expect potential negative outcomes.⁴ In the same manner, virtual violence may not be automatically wrong, simply because other users are hurt. Virtual violence that hurts other users should be considered morally problematic behavior, however, if it clearly violates (official or unwritten) rules of the virtual environment.

Virtual violence may not only be problematic, because it has detrimental effects on users (as agents) and sometimes also on other users. According to the reviewed literature, virtual violence may also be problematic, because it builds on the intention of a user to actually hurt another social agent. Intentions are meaningful. They tell something about the state of mind of agents. Intended behavior provides a message. If users intentionally engage in virtual violence, even though they intuitively perceive other virtual agents as social beings, their behavior would be problematic. Not because of negative effects, but because of an ill “wanting” of the users. Next to detrimental effects of virtual violence, it is probably the irritating symbolical significance of users’ intentions to do harm to others that raises moral concerns about virtual violence (e.g., Dorin 2004).

To conclude: Is virtual violence morally problematic behavior? It appears that virtual violence is an ethically relevant behavior. Virtual violence features problematic intentions and effects that are worth to be ethically considered and evaluated. In light of the literature that is reviewed in the present chapter, one may conclude that virtual violence is indeed morally problematic behavior.

References

- Anderson, C. A. & Bushman, B. J. (2002a). Human aggression. *Annual Review of Psychology*, 53, 27–51.
- Anderson, C. A. & Bushman, B. J. (2002b). The effects of media violence on society. *Science*, 295, 2377–2379.
- Anderson, C. A., Gentile, D. A., & Buckley, K. E. (2007). *Violent video game effects on children and adolescents*. New York: Oxford University Press.
- Anderson, C. A., Sakaoto, A., Gentile, et al. (2008). Longitudinal effects of violent video games on aggression in Japan and the United States. *Pediatrics*, 122, 1067–1072.
- Anderson, C.A. (2004). An update on the effects of violent video games. *Journal of Adolescence*, 27, 133–122.
- Anderson, M. & Anderson, S. L. (2007). Machine ethics: creating an ethical intelligent agent. *AI Magazine*, 28(4), 15–58.
- Bailenson, J., & Yee, N. (2005). Digital chameleons: Automatic assimilation of nonverbal gestures in immersive virtual environments. *Psychological Science*, 16, 814–819.
- Bargh, J. A. (1994). The Four Horsemen of automaticity: Awareness, efficiency, intention, and control in social cognition. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., pp. 1–40). Hillsdale, NJ: Erlbaum.
- Baron, A. & Richardson, DR (1994). *Human Aggression* (2nd ed.). New York: Plenum.
- Bartle, R. (2009, March). MMO morality. Invited keynote held at the Game Cultures Conference, March 19 – 21, Magdeburg, Germany.
- Berkowitz, L. (1993). *Aggression. Its causes, consequences, and control*. Philadelphia, PA: Temple University Press.
- Blascovich, J. & McCall, B. (in press). Attitudes in virtual reality. In J. Forgas, W. Crano & J. Cooper (Eds.), *Attitudes and persuasion*. Publisher [?]
- Boden, M. A. (2000). Autopoiesis and life. *Cognitive Science Quarterly*, 1, 117–145.
- Bredemeier, B. (1985). Moral reasoning and the perceived legitimacy of intentionally injurious sport acts. *Journal of Sport Psychology*, 7, 110–124.
- Bredemeier, B., & Shields, D. (1986). Athletic aggression: An issue of contextual morality. *Sociology of Sport Journal*, 3, 15–28.
- Bufacchi, V. (2004). Why is violence bad? *American Philosophical Quarterly*, 41 (2), 169–80.
- Bushman, B. J. & Anderson, C. A. (2001). Media violence and the American public. Scientific facts versus media misinformation. *American Psychologist*, 56(6/7), 477–489.
- Bushman, B. J. & Anderson, C. A. (2002). Violent video games and hostile expectations: A test of the general aggression model. *Personality and Social Psychology Bulletin*, 28(12), 1679–1686.
- Dorin, A. (2004). Building artificial life for play. *Artificial Life*, 10(1), 88–112.
- Eastin, M. S., & Griffiths, R. P. (2006). Beyond the shooter game: Examining presence and hostile outcomes among male game players. *Communication Research*, 33(6), 448–466.
- Elton, M. (2000). Should vegetarians play video games? *Philosophical Papers*, 29(1), 21–42.
- Elton, M. (2009). Robots and rights: The ethical demands of artificial agents. <http://www.abdn.ac.uk/philosophy/endsandmeans/vol1no2/elton.shtml>.
- Epstein, S., & Pacini, R. (1999). Some basic issues regarding dual-process theories from the perspective of cognitive-experiential self-theory. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 462–482). New York: Guilford.
- Floridi, L. & Sanders, J. W. (2004). *On the morality of artificial agents*. *Minds and Machine*, 14, 349–379.

- Geen, R. G. & Thomas, S. L. (1986). The immediate effects of media violence on behavior. *Journal of Social Issues*, 42, 7 – 27.
- Green, M. C., Brock, T. C., & Kaufman, G. F. (2004). Understanding media enjoyment: The role of transportation into narrative worlds. *Communication Theory*, 14, 311-327.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814-834.
- Harris, P.L. (2000). *The work of the imagination*. Oxford: Blackwell.
- Hartmann, T. & Vorderer, P. (in press). It's okay to shoot a character. Moral disengagement in violent video games. *Journal of Communication*.
- Hartmann, T., Toz, E., & Brandon, M. (in press). 'Just a game'? Unjust virtual violence produces guilt in empathetic players. *Media Psychology*.
- Hopf, W. H., Huber, G. L., & Weiss, R. H. (2008). Media violence, youth violence, and achievement in school - a two-year longitudinal study. *Journal of Media Psychology*, 20(3), 79-96.
- Huff, C., Miller, K. & Johnson, D. G. (2003). Virtual harms and real responsibility: a rape in cyberspace. *IEEE Technology and Society Magazine*, 23(2).
- Hunter, D. & Lastowka, F. G. (2003). To kill an avatar. Legal Affairs. Retrieved online from http://www.legalaffairs.org/issues/July-August-2003/feature_hunter_julaug03.htm.
- Klimmt, C., Schmid, H., Nosper, A., Hartmann, T. & Vorderer, P. (2006). How players manage moral concerns to make video game violence enjoyable. *Communications - the European Journal of Communication Research*, 31(3), 309-328.
- Konijn, E., Tanis, M., Utz, S. & Linden, A. (2008). Mediated interpersonal communication. Mahwah, NY: Lawrence Erlbaum Associates.
- Kugler, K., & Jones, W. H. (1992). On conceptualizing and assessing guilt. *Journal of Personality and Social Psychology*, 62(2), 318-327.
- Ladas, M. (2002). *Brutale Spiele(r)? Wirkung und Nutzung von Gewalt in Computerspielen* [Brutal players? Use and effects of violent video games]. Frankfurt / M.: Peter Lang.
- Luck, M. (2009). The gamer's dilemma: An analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia. *Ethics and Information Technology*, 11, 31–36.
- Mathiak, K & Weber, R. (2006). Towards brain correlates of natural behavior: fMRI during violent video games. *Human Brain Mapping*, 27, 948-956.
- McCall, C. Blascovich, J., Young, A. & Persky, S. (2009). Proxemic behaviors as predictors of aggression towards Black (but not White) males in an immersive virtual environment. *Social Influence*, 4(2), 138 – 154.
- McCormick, M. (2001). Is it wrong to play violent video games? *Ethics and Information Technology*, 3, 277-287.
- McDermott, D. (2007). Artificial intelligence and consciousness. In P. D. Zelazo, M. Moscovitch & E. Thompson (Eds.), *The Cambridge Handbook of Consciousness* (pp. 117-150). Cambridge: Cambridge University Press.
- Meehan, M., Insko, B., Whitton, M., & Brooks, F.P. (2002) Physiological measures of presence in stressful virtual environments, *ACM Transactions on Graphics, Proceedings of ACM SIGGRAPH*, 21(3), 645-653.
- Möller, I., & Krahé, B. (2009). Exposure to violent video games and aggression in German adolescents. *Aggressive Behavior*, 35, 75-89.
- Paik, H. & Comstock, G. (1994). The effects of television violence on antisocial behavior: A meta-analysis. *Communication Research*, 21, 516-546.
- Pizarro, D.A., Detweiler-Bedell, B., & Bloom, P. (2006). The creativity of everyday moral reasoning: Empathy, disgust and moral persuasion. In J. Kaufman & J. Baer (Eds.), *Creativity and Reason in Cognitive Development* (pp. 81 - 98). Cambridge University Press.
- Powers, T. M. (2003). Real wrongs in virtual communities. *Ethics and Information Technology*, 5, 191–198.

- Powers, T.M. (2003). Real wrongs in virtual communities. *Ethics and Information Technology*, 5, 191–198.
- Ray, T. S. (1992). An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer & S. Rasmussen (Eds.), *Artificial life II* (pp. 371–408). Redwood City, CA: Addison-Wesley.
- Roseman, I. J. (2001). A model of appraisal in the emotion system. In K. Scherer, A. Schorr & T. Johnstone (Eds.), *Appraisal processes in emotion. Theory, Methods, Research* (pp. 68 – 91). New York: Oxford University Press.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417–457.
- Searle, J. (1990, January). Is the brain's mind a computer program? *Scientific American* 262, 26 – 31.
- Shapiro, M. A. & McDonald, D. G. (1992). I'm not a real doctor, but I play one in virtual reality: Implications of virtual reality for judgments about reality. *Journal of Communication*, 42, 94–114.
- Slater, M.D., Henry, K.L., Swaim, R.C., & Anderson, L.L. (2003). Violent media content aggressiveness in adolescents. *Communication Research*, 30(6), 713–736.
- Smith, E. R., & DeCoster, J. (2000). Dual process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4, 108–131.
- Spafford, E. H. (1989). Computer viruses as artificial life. Manuscript retrieved online from <http://www.scs.carleton.ca/~soma/biosec/readings/spafford-viruses.pdf>.
- Suler, J.R. & Phillips, W. (1998). The bad boys of cyberspace: Deviant behavior in multimedia chat communities. *CyberPsychology and Behavior*, 1, 275–294.
- Weber, R., Ritterfeld, U., & Mathiak K. (2006). Does playing violent video games induce aggression? Empirical evidence of a functional magnetic resonance imaging study. *Media Psychology*, Vol. 8, No. 1, Pages 39–60.
- Whang, L. S-M. & Chang, G. (2004). Lifestyles of virtual world residents: Living in the on-line game “Lineage”. *CyberPsychology & Behavior*, 7(5), 592 – 600.
- Wolfendale, J. (2007). My avatar, my self: Virtual harm and attachment. *Ethics and Information Technology*, 9, 111–119.
- Wonderly, M. (2008). A Humean approach to assessing the moral significance of ultra-violent video games. *Ethics and Information Technology*, 10, 1–10.
- Yee, N., Bailenson, J.N., & Ducheneaut, N. (2009). The Proteus Effect: Implications of transformed digital self-representation on online and offline behavior. *Communication Research*, 36(2), 285–312.
- Zillmann, D. (2006). Dramaturgy for Emotions from Fictional Narration. In J. Bryant & P. Vorderer (Eds.), *Psychology of entertainment* (pp. 215–238). Mahwah, NJ: Lawrence Erlbaum Associates.

Notes

¹The intention-based approach to judge virtual violence will also be considered when tackling question 1. The underlying question is, if users indeed intend to do harm to other social beings (as presumed in the definition of virtual violence), if they play violent video games or commit violent acts in other virtual environments.

²It needs to be noted, however, that the conditions that constitute life or a living agent are lively debated among philosophers, biologists, computer scientists, and engineers (e.g., Ray 1992). Readers that like to learn more about this debate may also consult the scientific journal “Artificial Life”.

³It has to be noted that the main goal of research on artificial life is to *model* life in order to explore and understand it. The debate about *actually* living software systems (instead of software that merely imitates, simulates, or models life) is not of central importance to the general field.

⁴The given example may also be misleading, as it may not fully match the definition of violence that underlies the present chapter. An important part of the definition of violence is that the potential victim apparently tries to *avoid* the harm-doing. The opponents of a box-fight (or in any other violent sports competition) may certainly appear to avoid being harmed. However, the fact that they willingly entered the violent situation beforehand contradicts the notion of a victim that tries to avoid being harmed. The same applies to users that willingly entered a violent virtual environment. Accordingly, it may not be fully correct to label behavior “violent” that aims to do harm to those who willingly entered the violent situation (and expected the possibility to be harmed). This is not to say that virtual violence would never be harmful under those conditions, because it can still “amount to destruction of property, albeit in an environment where that is an expected outcome of participation.” (Dorin 2004, p. 108).