CHAPTER THREE

Virtual Entities, Environments, Worlds and Reality

Suggested Definitions and Taxonomy

Johnny Hartz Søraker

Introduction

In the current literature on 'virtuality' and related terms, there is hardly any consensus on what these terms should refer to, nor what their defining characteristics are. At one end of the spectrum, computer scientists often use the term to denote a method of simulating a piece of computer hardware by other means, for instance '*virtual* memory'. At the other end of the spectrum, we find science-fiction tales of humans immersed in environments that simulate the actual world in such richness and accuracy that the two cannot be distinguished. In between, we find more mundane uses of the term, as applied to computer games and Web sites. This conceptual muddle is unfortunate, especially when our goal is to evaluate their societal and ethical impact. For instance, if we ask the question of how far trust may be established in virtual environments, our answers will (or *should*) differ depending on whether we mean 'virtual environment' to refer to chat rooms, discussion forums, email, MMOs, text messaging, and so forth. The purpose of this paper is to distinguish between four related yet fundamentally different categories of virtuality: 'virtual' (generic), 'virtual environment', 'virtual world', and 'virtual reality' and to show how they give rise to different sets of considerations when evaluating their potential impact on their users.

First of all, as if referring to virtual entities is not difficult enough, the seemingly vanishing line between real and virtual prompts us to be precise about what is 'real'

as well. As I will return to below, many virtual states of affairs are just as real as (or no less illusory than) 'real reality'. Thus, I agree with Albert Borgmann's claim that 'a distinctive term for real reality needs to be found now that reality can be virtual' (Borgmann, 1999, p. 256).[1] Following his suggestion, I will use 'actual reality' when I refer to what is often misleadingly characterized as 'the real world'. I will sometimes use 'physical reality' when I explicitly address the physical as opposed to non-physical aspects of reality. As Borgmann does, I will sometimes use 'real' and 'the real thing' when there is no danger of confusion or when I address claims that specifically invoke these terms.

'Virtual' is often defined along the same lines as 'quasi-', 'pseudo-' or 'almost the same as'; something that is almost but not quite real, or something real without being *actual* (cf. Shields, 2003, p. 25, inspired by Proust and Deleuze). This usage is reflected in our daily language when we speak of something being *virtually* something else. In the same vein, Jaron Lanier, one of the early pioneers in virtual reality research who also popularized the term, explains that for something to be virtual 'it has to be indistinguishable [from the actual entity] in some practical context, while it remains distinguishable in another' (Lanier, 1999). Although this description gives us some idea what virtuality refers to and the notion of 'similar but not equal' will later prove important, the description is too imprecise and inclusive. First, it does not single out which properties of the entity in question are not actual. For instance, the term 'virtual soldier' can be applied to someone who is a human being and almost a soldier, or it can refer to 'someone' who is not a physical being, but has the properties of an actual soldier (e.g., a video game character). Second, if we admit anything that is 'almost but not quite real' in our definition, this seems to entail that vivid memories, dreams, movies and rainbows are virtual as well—which they are not, as I will return to shortly. As these examples show, a definition along the lines of 'almost real' fails to signify both how 'virtual' differs from 'actual' and, just as importantly, how it differs from thought constructs, dreams and illusions. Still, it is important to keep this aspect in our definition, because many of the philosophical issues arise precisely because our theoretical and ethical frameworks become difficult to apply when only *some* of a virtual entity's properties are actual. The challenge lies in saving the intuition that lies behind virtuality as something 'almost real', while avoiding the lack of precision that often follows.

My first step will be to define virtual as a generic term, but when I exclude certain criteria as necessary for such a broad conception of virtuality, I will return to some of these when outlining the different sub-categories. A promising starting point is to look at how 'virtual' relates to what is often referred to as 'new media'—and why it is so widespread to employ spatial metaphors.

Virtual as new media—the 'place' metaphor

There are two features that are often attributed to new media: instant, global communication, and digitized information. However, these properties are not unique to new media. First, different forms of telecommunication have been used for long-distance, instant communication since the early 20th century. Although computer-mediated communication further improves the availability, speed and reliability of long-distance communication, this does not in itself constitute a qualitative shift from traditional forms of media and telecommunication. Second, although everything virtual is ultimately digital in nature (cf. Figure 3-2), many forms of digital technology are not media and the term 'digital' is too broad to be of much value in defining new media, let alone to understand what is so special about it. This is also evidenced by the fact that many entities are digital without thereby being virtual, such as a digital radio or digital camera. Thus, a more promising way of differentiating between old and new media is to look at the fundamental way in which information transfer takes place in the respective forms.

In classical definitions of media, the information exchange is seen as a linear process where a *sender* sends a *message* through a *medium* to a *receiver*. For instance, in the influential transmission model of communication, originally developed by Shannon and Weaver (1949/1998), communication is defined as a message that originates from an *information source,* is encoded by a *transmitter,* sent through a *channel,* decoded by a *receiver* and ultimately arrives at a *destination*. In this model, as well as in daily language, the medium is regarded as a *channel* that we transfer information *through*. This is precisely what is subverted in new media. New media are not channels, but more akin to *places*. Virtual entities can be located in a geometric three-dimensional environment that literally resembles our traditional notion of a *geometric* space. However, they can also be located in a *topological* space, which is an abstract place where virtual entities are *located* and subjected to abstract ordering principles that define their topological location in relation to other virtual entities. For instance, a virtual document can be *within* a virtual folder and *moved* to another, you can go *back* and *forth* between Web sites, a hard drive can be *full* or *empty*, you can *up*load and *down*load files, and so forth (cf. Brey, 1998).

A similar but somewhat misleading view has been presented by Jonathan Steuer, who illustrates the difference between traditional and modern forms of *communication* in Figure 3-1.

As I will return to in more detail below, I think Steuer is mistaken in applying this model to all forms of modern communication and his inclusive use of the term 'virtual reality' blurs rather than clarifies what 'virtual' refers to. A regular conversation on the phone can be described metaphorically as being 'electronically pres-

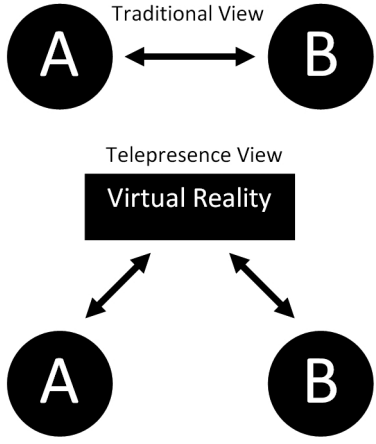

Figure 3-1: The 'traditional' versus the 'telepresence' view of communication (facsimile from Steuer 1992, p. 77)

ent in the same virtual reality' (Steuer, 1992, p. 78), but this metaphorical description is far removed from the way in which virtual worlds *literally* take place in a shared space. However, the model can be used to illustrate one of the crucial characteristics of all things virtual. With virtual environments, in a very literal sense, information is better described as being *made available* rather than *transferred*. As such, the focus in new media, when properly nuanced, is indeed on 'the relationship between an individual who is both a sender and receiver, and on the mediated environment with which he or she interacts' (Steuer, 1992, p. 78). In other words, the medium is not a channel through which we transfer information, but more like a place in which the information remains for a significant duration of time. At the same time, there is no clear distinction between the medium and the information. For instance, if we simply extract the text from a Website or virtual environment, much of the information is lost, including the way it is presented, its topological or geometric relation to other information, and the ways in which we can interact with it. Thus, with virtuality, far beyond a mere metaphorical sense, the message is the medium; the environment constitutes an important part of the information. Strictly speaking, it is not correct to say that these places do not physically exist *anywhere*, since they have to exist in form of a digital representation. However, they do not exist *as meaningful information in* the physical world, let alone as spatio-temporal objects. Indeed, the oft-mentioned oxymoronic features of virtuality arise from the fact that these environments, which constitute *and* contain information, seemingly do not exist in our physical world and can only be accessed using special equipment. This, in turn, is what prompts the familiar spatial metaphors such as cyber*space*, virtual *worlds* and info*sphere*.[2]

Among other things, this model shows how new media can offer many-to-many communication. In many-to-many communication, diverse information cannot be transferred simultaneously from many to many. Instead, a venue must be created in which many can leave information and many can retrieve information; anything else would amount to a cacophony of information overload. This model also conveys how new media can and must be *interactive*, because the same user can be both sender and receiver, and receive available information in accordance with the actions she carries out. Again, manipulating and choosing information (and the way it is presented) in this sense requires that the information is made available rather than transferred through channels. Interactivity entails that you do not choose between 'on' or 'off', which is the primary choice in the earlier sender-receiver model of communication; it entails choosing what, when and how (I will explain this in more detail below).

A review of suggested definitions

Although seeing new media and the virtual as a 'place' rather than a 'channel' brings out some of its important characteristics, this is also in many ways where the seemingly oxymoronic features stem from, and it does not give us a much clearer understanding of what it is or what is so special about it. A number of philosophers and computer scientists have struggled to propose clear and precise definitions, but there is little if any consensus. Before dealing with this problem myself, I will review a selection of promising definitions, and try to extract the criteria that will form the basis of the definitions I will propose towards the end of this paper.[3]

Stanovsky—computer-generated simulation and interactivity

Derek Stanovsky singles out 'computer-generated simulations' and 'interactivity' as the defining terms of virtuality. Note that the qualifier 'computer-generated' has been added to the simulation-criterion, thereby avoiding the criticism I raised against the overinclusiveness of 'simulation'. These, Stanovsky states, *might* be supplemented by three other qualities: being shared by multiple users, providing fully realistic sensory immersion, and enabling users to communicate and act over great distances (Stanovsky, 2004, p. 168). Since Stanovsky explicitly states that the latter are not necessary, he seems to imply that 'computer-generated simulations' and 'interactivity' *are* necessary. Although Stanovsky does not arrive at a concrete definition, his analysis accurately points out what I take to be the two central points of discussion if we are to get a better understanding of what virtuality is. First, in what

sense can actual entities be simulated or reproduced in virtual form and what is so special about *computer-generated* as opposed to other forms of simulation? Second, what is so special about the kind of human-computer *interactivity* that is facilitated by virtuality? Without a further elaboration of these points, which I will return to below, one runs the risk of regarding virtual experiences as being 'on more or less equal footing with the more usual forms of experience', or that "all experience is essentially virtual" (Stanovsky, 2004, pp. 172–173).

Brey—three-dimensionality and first-person perspective

Philip Brey defines virtual reality as 'a three-dimensional interactive computer-generated environment that incorporates a first-person perspective' (Brey, 1999, p. 5, 2008) and underlines that full immersion is not to be regarded as a necessary property. Part of the reason for excluding this criterion is that the range of technologies that encompass full immersion is still reserved for very few users. Hence, an ethics dedicated to those systems alone would hardly be worth the effort. He further qualifies his definition by stating that interactivity includes common peripherals such as mouse or gamepad, and that stereo vision, i.e., the actual sensation of being in a 3-D environment, is not a necessary requirement either. Thereby, Brey includes three-dimensional virtual environments that are projected onto a two dimensional computer monitor. Even if projected in this manner, Brey argues that it is essential that the virtual environment itself is three-dimensional, since it fosters an entirely different feeling of immersion and range of life-like actions.

Closely related, it is interesting to note that Brey prefers the concept 'first-person perspective' to 'immersion'. As mentioned above, immersion is a problematic concept since it is primarily a subjective concept—different people become immersed in different things to different degrees. Substituting this concept with 'first person perspective' avoids this problem. The reason why Brey includes the first-person requirement is that it 'suggests a degree of immersion in a world, rather than the experience of a world that can be (partially) controlled from the outside' (Brey, 1999, p. 6). The kind of immersion Brey attributes to having a first-person perspective is clearly important, but the term itself is slightly ambivalent. First, most forms of screen-based VR offer the opportunity to choose freely one's point of view; whether or not one sees the world through the eyes of a virtual body—an *avatar*—does not seem to make much difference. Seeing the world through the eyes of an avatar only makes a profound difference if the perspective replaces our real-world view, e.g., by using a head-mounted display or special VR goggles. However, on this strict interpretation of first-person perspective, which I will later refer to as a first person *view*, many of the applications Brey counts as instances of VR would be excluded. Thus, the reason for using the first-person perspective will

prove important in our definition, but due to the ambiguity of the term, I will later suggest the term 'indexicality' instead.

Sherman and Craig

Finally, Sherman and Craig offer the following definition in their *Understanding Virtual Reality* (2003):

> …a medium composed of interactive computer simulations that sense the participant's position and actions and replace or augment the feedback to one or more senses, giving the feeling of being mentally immersed in the simulation (Sherman & Craig, 2003, p. 13).

Although this definition provides further insight into the essence of virtuality, its meaning differs dramatically depending on how we interpret 'participant'. If, by 'participant's position', Sherman and Craig refer to the position of our physical bodies, then the definition excludes any kind of virtuality in which actions are mediated by an avatar. On this reading, virtuality refers only to those technologies that employ a body suit or other means of motion tracking. Admittedly, this is a distinct and important characteristic of what I will refer to as the subcategory virtual *reality*, but it excludes many of the most interesting and widely available forms of virtuality, in particular any form of virtual environment experienced through a computer monitor. Thus, Sherman and Craig's definition is more apt to define the high-tech forms of virtual *reality*, such as the CAVE,[4] rather than virtuality in general. Sherman and Craig do strike at the core of what virtuality is about in describing it as 'a medium composed of interactive computer simulations' and I will shortly return to their important emphasis on the active role of the computer simulation.

'Virtual' as interactive computer simulation

It is interesting to note that all the definitions discussed include some form of 'computer simulation' and 'interactivity'. I will argue that these are indeed the defining criteria for a generic definition of 'virtual', but, as with the other criteria reviewed, none of these are sufficient on their own. There are of course entities that are interactive without being computer simulations, for instance a car, and entities that are computer simulations without being interactive, for instance when a calculation is performed on the basis of pre-programmed parameters without human intervention. The combination of the two however—interactive computer simulations—will provide the foundation of my definitions and taxonomy of virtuality. Based on pre-

vious remarks on the over-inclusiveness of mere 'simulation' and the unique, *interactive* characteristics of virtuality as a new medium, this is a promising starting point for finding a lowest common denominator of all things virtual. 'Simulation' is also not a precise description of virtuality on its own, since it includes a number of situations that cannot readily be described as virtual. In other words, the unique epistemological and ontological status of virtuality disappears if we analyse the phenomenon on a par with, say, an architect's use of miniature models. At the same time, 'simulation' is a more precise term than 'almost the same as' while retaining the notion of something virtual as actual in *some* respects but not in others (as reflected in Lanier's definition above). 'Simulation' accurately captures this property because it entails an analysis of x in order to infer to y. That is, for an x to be a simulation of y, this requires that x cannot be exactly the same as y *and* that x must have some relevant properties in common with y. A miniature airplane will not be a simulation of an actual airplane if the two are the same, nor if the miniature airplane does not have any relevant properties in common with the actual airplane.

To clarify in what sense virtual entities and environments are essentially simulated in nature, a few clarifications are in order before I turn to how additional features constitute different subcategories of virtuality. First, what are they simulations *of*?

Computer simulation of types and tokens

It might sound counterintuitive to define anything virtual as an interactive computer simulation, since this entails that all virtual entities are simulations *of* something. How does this square with the fact that virtual entities do not always correspond to actual entities? After all, we can have computer simulations of all kinds of nonexisting entities, such as fantasy realms and fantasy creatures. This objection is valid, but only if we presuppose a notion of simulation in which the simulated entity must exist as a physical entity. On such a presupposition, a virtual house would fit the definition, whereas a virtual zombie would not. Thus, the notion of 'virtual' as 'simulation of actual' requires some important clarifications. On my usage of the term, 'simulation' can range from token-token to type-type relationships—and, slightly more controversially, it can include simulation of both existing (concrete) and nonexisting (abstract) entities.

First, the distinction between types and tokens is a distinction between a general kind of thing (e.g., horses) and its particular instances ('Clever Hans' or some other particular horse). For instance, we can talk of the horse *type* being thousands of years old despite the fact that no horse *token* would live to that age. There are primarily two competing views on the relation between types and tokens (cf. Wetzel, 2007). On a Platonic view, types (or 'forms') exist independently of their tokens, if any, and do not exist anywhere in space-time. On an Aristotelian view, the types cannot exist

independently from their tokens. According to Aristotle, if all tokens of a type (all physically existing horses) were to cease to exist, then the type ('horse') would cease to exist as well. According to Plato, whether the token exists or not does not determine whether the type ('the Idea') exists. What is interesting about these differing views is that from a Platonic perspective, there is no fundamental difference between the type 'horse' and the type 'pegasi', because whether or not they exist is irrelevant. From an Aristotelian perspective, however, this presents a curious dilemma. What kind of concept is 'pegasi' if it has no physical instances? For present purposes, this problem prompts a second distinction between 'abstract' and 'concrete'.

The distinction between abstract and concrete is equally controversial, and a number of different interpretations have been put forward (cf. Rosen, 2006). According to what Lewis calls the 'way of negation' (Lewis, 1986), 'abstract' is typically defined in terms of what it is not—especially which physical properties that are lacking. Perhaps the most promising definition along these lines is that abstract tokens are distinguished by the fact that they have no spatiotemporal location and (hence) no mechanico-causal powers. For instance, a physical horse exists in spacetime and can bring about all kinds of changes in the physical world, whereas an abstract Pegasus does not exist in space and can bring about no changes in the physical world. On this definition, 'concrete' refers to any (type of) entity that occupies physical space and is, consequently, capable of causing mechanico-causal effects in the physical world. 'Abstract' refers to all entities that are not concrete in this manner, be they types (e.g., centaurs) or tokens (e.g., the centaur Chiron).

If we follow the distinctions between concrete/abstract and type/token, we get four different kinds of simulation:

> **Concrete type simulations**: Simulation of types that are instantiated in concrete tokens with determinate spatiotemporal location and causal powers (virtual horses, virtual mountains etc.)
>
> **Abstract type simulations**: Simulation of types that are not instantiated in concrete tokens, thus have no determinate spatiotemporal location and causal powers (virtual pegasi, virtual fantasy realms etc.)
>
> **Concrete token simulations**: Simulation of concrete tokens that have a determinate spatiotemporal location and causal powers (virtual 'Clever Hans', virtual Mount Everest etc.)
>
> **Abstract token simulations**: Simulation of abstract tokens that have no determinate spatiotemporal location and causal powers (virtual Pegasus [son of Poseidon], virtual Valhalla etc.)

A fourfold table might clarify further:

	CONCRETE	ABSTRACT
TYPE	Horses, Mountains	Pegasi, Fantasy realms
TOKEN	'Clever Hans', Mount Everest	Pegasus, Valhalla

Table 3-1: Fourfold diagram of types/tokens and concrete/abstract

The important point is that any entity or environment that is virtual is necessarily computer-simulated in one of the senses above. For instance, a virtual library is an interactive concrete *type* simulation (a simulation of something that does exist, but not of one particular entity) whereas the virtual Amsterdam library is an interactive concrete *token* simulation (a simulation of something that does exist, and which is a particular entity, i.e., the actual Amsterdam library). A virtual pegasus is an interactive, abstract *type* simulation (a simulation of something that does not exist and which is not a particular entity) whereas a virtual Pegasus (the token son of Poseidon) is an interactive, abstract *token* simulation (a simulation of something that does not exist, but which is a particular entity). I will later use these distinctions to indicate different forms of virtual-actual relations.

Finally, it could also be objected to my definition that some of the things we refer to as virtual are not in themselves computer-simulated and interactive. For instance, a virtual community is not, strictly speaking, a computer-simulated, interactive community. However, a virtual community is virtual in the sense that it is *made possible by* an interactive computer simulation—in the same sense as an actual community is made possible by the shared venues, shared community objects and so forth. Indeed, one of the essential characteristics of these kinds of social constructed entities, institutional entities in particular, is that their nature is only partly, if at all, determined by their physical properties (Searle, 1995). More generally, in much the same manner that many actual entities and events require an explanation beyond being made possible by physical matter, many virtual entities require an explanation beyond being made possible by an interactive computer simulation.

The computational underpinning

A significant degree of regularity is of fundamental importance to the epistemological status of virtual entities and is secured by the role of the computer that underpins the virtual worlds. This can perhaps best be described by way of a

comparison with different philosophical accounts of regularity in the physical world. The similarities and differences between virtual and physical regularity can be illustrated by way of the age-old philosophical discussion of how causality is possible in the physical world. David Hume famously argued that our notion of causality is based on nothing but events that closely succeed each other and therefore become associated into what we refer to as cause and effect:

> Suppose two objects to be presented to us, of which the one is the cause and the other the effect; it is plain, that from the simple consideration of one, or both these objects we never shall perceive the tie by which they are united, or be able certainly to pronounce, that there is a connexion betwixt them (Hume, 1739/2008)

If we see a billiard ball collide with another ball, we cannot experience the supposed causality between them. In other words, there is no magical application of a physical law that determines that certain causes have certain effects. Nicolas Malebranche (1688/1997) struggles with the same problem and finds himself having to postulate an omnipotent God that must actively intervene and make sure that causes are followed by regular effects. Whatever the correct account of causality in the real world might be, Hume and Malebranche point to an important characteristic of virtual worlds. If two billiard balls collide with each other in a virtual world, there *is* a connection between them, namely the computer event that is triggered on that specific occasion and determines what the effect of that cause is going to be. That is, the computer fulfils more or less the same function as Malebranche's God, by constantly having to intervene on every occasion where a particular effect should follow a cause. In more technical terms, virtual environments are characterized by the fact that there is a causal engine that determines the effects of certain causes, for instance by altering the movement of one ball when hit by another on the basis of the mathematical properties of the balls (velocity, direction, impact angle, mass etc.).

This unique characteristic of computer-generated simulation is precisely what makes the virtual fundamentally different from mere products of the mind and at the same time fundamentally different from the real world. The equivalent of Malebranche's God and whatever it is that Hume is missing when we experience an effect following a cause in the physical world can be found within the computer technology that underpins a computer simulation. Although computer-simulated laws of physics are different from laws of physics (e.g., more irregular and subject to lack of accuracy), they can still add a high degree of regularity. That is, the computer simulation can lend stability, persistence, predictability, and intersubjective availability to virtual events and entities, thereby separating the virtual from mere dreams and hallucinations. In more philosophical terms, the computer simulation

adds what Kant (1781/1997; KdRV) referred to as congruence between experiences.[5] That is, events do not happen at random, but are results of long, complicated and interconnected chains of events—and these events are subject to certain regularities that allow us to predict and explain the consequences of events and actions.

Most importantly, this entails that facts about virtual worlds can be epistemologically objective. If I claim that my virtual house is located on top of a virtual mountain on a specific island in *Second Life*, this is either true or false for two interrelated reasons. First, anyone can check whether this is the case or not (it is intersubjectively available) and second, the epistemological objectivity of these kinds of claims is made possible by the fact that the computer simulation gives rise to a regularity that allows for both accurate verification/falsification and prediction. This is also the reason why virtual environments are apt for scientific modelling, since the regularities can be created so as to resemble physical laws. Ultimately, it is this regularity that makes virtual acts worth undertaking, since they could not bring about desired consequences if there were no regularity and predictability. Indeed, it is this kind of regularity that makes virtual environments into something more than a spectacle—something more than the irregularity of dreams, illusions and hallucinations. This is in line with Kant's view that certain organizing principles (time, space, causality, permanence etc.) are necessary 'conditions of possibility' to make our experiences intelligible.

Thus, the reason why it is so important to include the computer in our definition of virtuality is to emphasize this point. The 'ultimate reality' of a virtual entity is rooted in the physical world, but only in a form that is not directly accessible (e.g., as strings of binary digits). Thus, truth claims about virtual entities are very similar to claims in general about the physical world. Our claims are usually not about whatever ultimate reality lies behind the world as it appears to us, but they are precisely about how that ultimate reality manifests itself in the world as it appears to us. For instance, it is still somewhat of a mystery what a gravitational force *really* is, but this does not prevent us from making truth claims and predictions on the basis of how the ultimate reality behind gravity manifests itself in the world as it appears to us. Just as we sometimes need instruments to assess the veracity of such claims, we need 'instruments' (screens, networks, peripherals etc.) to assess the veracity of claims about virtual worlds.

In short, it is the computer that underpins the virtual world that facilitates its epistemological objectivity and, in that respect, makes virtual worlds *similar* to the physical world and dramatically *different* from dreams, hallucinations and other products of the mind. At the same time, there are clearly differences between the virtual and the physical, precisely because a computer simulation is a necessary condition for the existence of the virtual world—and because the computer, in contrast with the laws of physics or Malebranche's God, is restricted by the available computational resources.

Interactivity

Although the importance of the computational underpinning captures some of the essence of virtuality, ranging from Websites to immersive VR, it is still limited in the sense that if these technologies were to provide nothing but passive perception, many of the interesting aspects of virtuality would be lost. It is certainly fascinating to watch three-dimensional computer-generated hologram artworks, but as long as they cannot be interacted with, they do not raise the same kind of ontological, epistemological and ethical puzzles. As Dominic McIver Lopes puts it: 'If virtual reality offers anything new it is the possibility for interaction with the occupants and furniture of the computer-generated environment', which is made possible 'precisely because of the special capabilities of computing technology' (Lopes, 2004, p. 110).

Most obviously, the ability to interact with the computer allows humans to communicate and interact with each other over great distances *through* the computer. Thus, what I will refer to as a 'virtual world' differs enormously from a hologram because of the ability to share the experience with others from across the globe. Although 'networked communication' is not necessary for something to be virtual, as I will return to below, it is an important part of virtual worlds and virtual communities, and *interactivity* is necessary for networked communication. Another obvious advantage of interactivity is that it is required in order to have an immersive three-dimensional, or what I will refer to as an indexical, experience. Simply put, you cannot 'look around' if the computer simulation has no idea where you or your virtual representation (the avatar) is looking, which requires interaction—either indirectly by some kind of steering mechanism or directly by means of motion detectors that report the direction of your head/eyes to the computer simulation. More fundamentally, some level of interactivity is necessary for all kinds of virtuality. In computer science terminology, an interactive computer simulation is sometimes aptly described as a computer simulation with a 'human-in-the-loop'.[6] Without a human in the loop, a computer simulation (running on a finite state computer) will be entirely deterministic in the sense that it will only do whatever it has been programmed to do—determined solely by pre-programmed variables. If the same variables are used, the computer simulation (running on a deterministic computer) will produce exactly the same result on subsequent runs.

Jonathan Steuer defines interactivity in terms of 'the extent to which users can participate in modifying the form and content of a mediated environment in real time' (Steuer, 1992, p. 84).[7] Defining human-computer interaction as the ability to modify the form and content of a mediated environment can be seen as setting the threshold too high, since a number of technologies that are sometimes referred to as virtual do not meet the requirements. Although I do in general strive

to define concepts according to common usage, this is a case where the term virtuality *ought* to be restricted. Some of the technologies hailed as interactive are in fact not. For instance, most forms of streamed media on the Internet are, in all important respects, the same as traditional forms of media if the user cannot, in a non-trivial manner, choose what to watch, when to watch it, and in which form to watch it—to, as it were, seek out the place where the content exists and interact with it. If we regard interactivity as a necessary requirement for virtuality, it follows that online, streaming, live media, for instance, do not qualify. Thus, terms like 'virtual TV channel' or 'virtual radio' are misnomers if the only thing that separates them from traditional media is that they are transmitted in digital rather than analogue form. The proper terms would be 'digital TV channel' or 'digital radio'. If we were to include these media in our definition of virtual, we would be left with a definition along the lines of 'not physical', or worse, 'something other than the old' since a 'real' TV channel is not physical either. Indeed, I believe the way in which 'virtual' has become a prefix for any kind of digitized version of traditional media and other phenomena is part of the reason for its present vagueness and ambiguity.

That being said, 'interactivity' should not be defined too strictly. In what I will shortly define as virtual environments (which includes, e.g., three-dimensional computer games), it is important that the environment is what Dominic McIver Lopes characterizes as *strongly* interactive. Strong interactivity entails that the 'users' inputs help determine the subsequent state of play' (Lopes, 2001, p. 68). In short, the state of affairs in a virtual environment should be largely determined by the users' actions. At the very least, this pertains to the location, orientation, abilities and so forth of your avatar, but the interactivity will of course be particularly strong when the state of the environment is fundamentally molded by the actions of the user. This is primarily the case in what I will refer to as persistent, multi-user virtual *worlds* where the state of the environment is a result of many years of user input from multiple users. When it comes to virtual *entities,* we should allow for a somewhat weaker sense of interactivity. This can be illustrated by way of game researcher Espen Aarseth's distinction between fictional and simulated objects (Aarseth, 2006). In the classic game *Return to Castle Wolfenstein,* the doors in the game are sometimes interactive (they can be opened, closed and so forth) but other doors are simply textures on the wall—and not interactive in a strong sense of the word. However, these doors may still be referred to as being computer-simulated and interactive, in the sense that they will stop your avatar from walking through them. The interactivity lies in the fact that the computer simulation must actively prevent your avatar from walking through the door. Hence, there is an interactive process between your avatar-mediated actions and the computer simulation that underpins the virtual environment.

'Virtual' defined

To summarize so far, we can define 'virtuality', and the adjective 'virtual' as interactive computer simulations. The definition consists of three elements:

- **'Computer':** The ontologically objective, physical 'grounding' that facilitates regularity and intersubjective availability, hence epistemological objectivity.
- **'Simulation':** entails that virtual entities are similar to actual entities in some respects and different in other (partly determined by whether they are abstract or concrete, types or tokens).
- **'Interactivity':** facilitates networked computing, immersion (indexicality) and virtual acts as modification of the form and content of a virtual environment.

Virtuality in this generic sense raises many philosophical problems on its own, primarily theoretical ones. However, the generic definition hides the fact that virtuality comes in many different forms, each presenting unique philosophical issues. In what follows, I will outline what I regard to be the three most important kinds of virtuality: virtual *environments,* virtual *worlds,* and virtual *reality*—which correspond to the criteria 'indexicality', 'multi-access' and 'first-person view'.

Virtual environments, indexicality and avatars

Recall that Brey singled out three-dimensionality and having a first-person perspective as a requirement for virtual reality, due to their ability to foster a form of immersion in the virtual environment. Although I agree with the reason for including these criteria, and find them more precise than the more commonly used 'immersion', the notion 'first-person perspective' is slightly misleading if defined loosely and too exclusive if defined strictly. First, if we define first-person perspective strictly (which I will later refer to as first-person *view)*, it requires one to have the ability to literally look around a virtual environment and change perspective by moving your head and/or eyes as you would in the physical world. Such a conception would exclude many of the applications that Brey includes, in particular all non-immersive (screen-based) forms of VR. If we define the term more loosely, however, the term becomes slightly misleading. Whether or not you *perceive* yourself as acting from a first- or third-person perspective in a screen-based form of VR does not make much difference. What makes a difference lies in the fact that a virtual environment can be 'interacted with from a single locus' (Brey, 1999, p. 6).

Interacting from a single location in the virtual environment without having a first-person *view* requires some kind of virtual entity that you control and through which you interact with other virtual entities. This is known as an avatar, which is the representation of the human user in a virtual environment, either in the form of a graphical object (usually in human-like form) or in the form of a nickname in non-graphical environments such as discussion forums and chat rooms.[8] In virtual worlds such as *Second Life* all interactions are carried out as if done from the location of your avatar, but you are free to choose whether you perceive the world from your avatar's eyes, from behind your avatar's shoulders, or from any 'God's eye' point of view. What is important is that the agent interacting with the environment is itself located in the environment.[9] Using an avatar that is located in a three-dimensional environment is important because it means that 1) the movements and actions of your avatar can be restricted in many of the same ways as in real life; 2) your avatar can be interacted with by the environment and, if in a multi-access environment (see below), seen and interacted with by others; 3) you can often configure your avatar to look however you wish thereby possibly creating stronger emotional ties between yourself and your representation; and 4) your avatar can engage in bodily acts that are similar to real life actions. These points, especially the latter, are responsible for many of the ethical problems that can arise in virtual environments, as for instance when users commit acts with their avatar that are unethical and/or illegal in the real world, such as virtual violence and virtual paedophilia.

Due to the ambiguity of terms like 'immersion' and 'first-person perspective', and the important differences caused by being an agent *in* a virtual environment, I propose 'indexical' as an alternative criterion. Admittedly, the term does not make immediate sense, but I have chosen it for two reasons. First, the term does not suffer from the ambiguity caused by the metaphorical use of terms like 'immersion' and 'first-person perspective' and I will reserve the notion of first-person *view*, defined in a strictly non-metaphorical manner, as the criterion for fully immersive virtual environments only—which I will refer to as virtual *reality* (see below).

Second, philosophers such as Ernst Tugendhat (1976) and Truls Wyller (2000) have given a precise meaning to the term that perfectly captures the importance of acting from a single locus. In philosophy, 'indexicality' is used in two different but related ways. In philosophy of language (and linguistics), indexicality refers to words whose meaning depends on the context in which they are uttered. For instance, saying that there is a tree to my left does not provide any knowledge about the location of the tree unless you know my spatio-temporal location and orientation. This means that terms such as 'there', 'here', 'to my left' and 'above me' do not make literal sense unless the person who uses the words has a known location and orientation in a three-dimensional space. This is what has spurred a second use of the term 'indexicality', inspired by neo-Kantian philosophers such as Tugendhat

and Wyller. In this second use, indexicality is not only a property of words, but also a property of our relation to our surroundings. As an example, trying to orient oneself by use of a map is entirely pointless if one does not know one's location on the map. Thus, stationary maps typically have a 'you are here' marker, which serves as an *index* from which you can orient yourself. In this sense of the word, 'indexicality' means to have a discreet, subjective (or egocentric) position from which we act and orient ourselves in a three-dimensional world (cf. Wyller, 2000, p. 39). In this form, the notion 'indexicality' captures both the three-dimensionality and the first-person criteria, and does not carry the metaphorical connotations implicit in terms like 'immersion'. When coupled with the interactivity criterion, which is a necessary criterion for all things virtual, indexicality also means that one is an *agent* at a specific place and, as such, *is present* at that place in one form or another—typically as an avatar. Thus, a first subcategory of virtuality is that which is not only computer-simulated and interactive, but also *indexical* as defined above. Since indexicality requires a three-dimensional space, I will refer to such an interactive and indexical computer simulation as a *virtual environment.*[10]

The reason I do not include indexicality as a criterion for virtuality in general is that terms such as virtual libraries, virtual banks and virtual universities make perfect sense even if they are not three-dimensional. Indeed, in virtue of being what Brey (2003) refers to as 'ontological reproductions', these are among the most important phenomena when it comes to the impact of virtual entities on our daily lives. These are the types of entities that have been socially constructed (a library, bank or university would not be what they are if there were no conscious beings), which means that their existence is not closely tied to particular physical properties. The same type of social reality can also be constructed from within what I will refer to as virtual *worlds,* as made possible by the presence of *multiple users* within a virtual environment.

Virtual worlds, multi-access and virtual communities

Virtual environments, as defined above, have been prevalent since the 1980s, in the form of single-player computer games. The addition of multiple access in virtual environments only became widely popular with the introduction of affordable broadband in the late 1990s. Text-based virtual communities, however, have been relatively widespread since the 1980's—famous examples including the WELL (Whole Earth 'Lectronic Link) and LambdaMOO. By 'multiple access' (or the 'multi-access-criterion') I mean the simultaneous presence of multiple participants in a computer-simulated interactive environment (three-dimensional or otherwise), where the participants can communicate with each other. In an interactive computer simulation,

this requires (graphical) representations of the participants and the use of networked communication.[11] It also requires (or, facilitates) a 'persistent space', meaning that you normally cannot *pause* a virtual world. Indeed, the reason I have chosen 'multiple access', rather than 'multiple users' or similar, is that access captures the importance of there being a persistent world that is persistently 'there' for users to access—a world embodying objects that continue to exist, and events that transpire also when you are logged out. Virtual worlds exist for users to go in and out of, and your actions and their consequences take place in the same space as the actions and consequences of other users. Thus, multi-access, and hence persistent space, makes the place metaphor even more apt; the 'place' exists whether you are there or not, and not only in a topological meaning. This also explains why virtual worlds have to be real-time, since they must allow one user's actions to be immediately visible to other users. The multi-access criterion is difficult because it cuts across very differing forms of virtuality, and is absent even in some of the most technologically advanced forms. There is a profound difference between multiple users that, on the one hand, meet and communicate in a virtual world proper, and on the other, multiple users that meet and communicate in a non-graphical discussion forum or chat room. What they have in common is that users communicate through a representation of themselves, thus allowing for pseudonymous communication. However, the philosophical and ethical challenges become much more complicated with the introduction of a *graphical* representation, since this allows for re-creation of acts, events and experiences that cannot be carried out in the actual world, be it for physical, economical or ethical reasons. Hence, I will use the established term 'virtual community' to describe communication between multiple users that is computer mediated and interactive, but not necessarily taking place in a three-dimensional environment. Because of the significant difference between virtual communities on the Web and virtual communities inside virtual *environments*, I propose that computer-simulated, interactive, indexical, *multi-user* environments should be referred to as *virtual worlds*.[12] Among the reasons for choosing 'world', in contrast with environment, is that 'world' is often used to signify a collection of people, or the earth *as inhabited*—hence expressions like 'the modern world' and 'the ancient world', which implies different worlds on the same earth.[13]

I have defined virtual worlds as interactive, computer-simulated, indexical, multi-access environments, which means that all virtual worlds are graphical, three-dimensional, persistent environments that allow for user-to-user interaction. Although these features are shared by all virtual worlds, it is important to note that there are still a number of differences between the many hundred different virtual worlds out there.[14] Indeed, Sarah Robbins-Bell (2008) is developing a faceted classification scheme of virtual worlds, according to which virtual worlds are classified according to such categories as physics, type of objectives, dominant form of communication, object ownership, avatar customizability, access model, user-to-user

relationship and formation of communities. Although *Second Life* and *World of Warcraft* are both virtual worlds according to my definition, I agree with Robbins-Bell's elaboration of the nuances within this category, according to which *Second Life* and *World of Warcraft* differ in almost all respects.

So far, I have defined 'virtual' as something that is computer-simulated and interactive, and 'virtual environment' as a subcategory that requires indexicality in addition. 'Virtual worlds' are, in turn, a subcategory of 'virtual environments', further requiring multiple access as outlined above. I now turn to 'virtual reality', which is also a sub-category of 'virtual *environment*' in that it requires indexicality. However, 'virtual reality' does not require multi-access, but instead what I will refer to as a genuine 'first-person view'.

Virtual reality and first-person view

For virtual reality, indexicality is necessary but not sufficient; it requires your indexical location, orientation and movements in the virtual environment to actually correspond to the indexical location, orientation and movement of your physical body—it requires what I will refer to as a genuine 'first-person view'.

Having a first-person view, in a strict sense, means that what you perceive is determined by the location and orientation of your eyes. For example, a first-person view restricts you from seeing an object from different angles without either rotating the object or moving your eyes and body. This strict definition entails that, in order to have a first-person view in a virtual environment, 1) you must be in a three-dimensional environment, 2) you must physically move (a part of) your body in order to look around and navigate, and 3) your movements and location in the virtual environment corresponds to your movements and location in the actual world. That is, there is a 'concrete token' relationship between your movements and location in the virtual reality and the movements and location of (parts of) your physical body. Thereby, 'first-person view' encompasses all the essential characteristics of high-end forms of virtual reality and renders indexicality as redundant for the definition of virtual reality. This is reflected in the fact that all high-end virtual reality systems require a device that directly feeds the virtual environment to your eyes, by means of either a head-mounted display or VR goggles.[15] When using transparent goggles that convert a stereographic two-dimensional image into the appearance of a three-dimensional environment, you move and look around just as you would in real life. Head-mounted displays, however, require motion tracking in order to adjust the participant's field of view according to the movements and orientation of the head. Having a first-person view in this sense truly makes the virtual experience lifelike. This is reflected in the fact that virtual reality first-person view can induce,

and is often used to treat, phobias such as fear of heights or arachnophobia (see e.g., North, North, & Coble, 1998)

Based on the considerations above, I take first-person view to be the minimal requirement of virtual *reality,* since it truly *immerses* a participant in the virtual reality, far beyond any metaphorical sense. The importance and necessity of the connection between movement of the eyes and spontaneously altered perception is nicely captured by Edmund Husserl:

> If the eye turns in a certain way, then so does the 'image'; if it turns differently in some definite fashion, then so does the image alter differently, in correspondence. We constantly find here this two-fold articulation…Perception is without exception a *unitary accomplishment* which arises essentially out of the playing together of two *correlatively related functions*…In virtue of [these processes] one and the same external world is present to me. (Husserl, 1989, p. 63 [58], italics in original)

An interesting way of illustrating Husserl's phrase 'one and the same external world' is that with the term virtual *reality*, in contrast with virtual environments and virtual worlds, we intuitively resist its plural form. The reason is precisely that being in a reality requires having a first-person view, and it is impossible to have more than one first-person view. That is, you cannot participate in multiple virtual realities simultaneously any more than you can be in more than one spatiotemporal place simultaneously in physical reality. You can have as many avatars in as many virtual worlds or environments as you like, but as soon as you experience the virtual reality 'through your own eyes' you can only be in one virtual reality at a time.

The virtual reality can be augmented by the addition of more comprehensive forms of motion tracking and sensory feedback. By using highly advanced data gloves, one can have tactile perception, for instance by feeling resistance when one grabs an object.[16] It should be noted that motion tracking comes in widely differing degrees. The most advanced forms of VR can implement so-called 'full body immersion' by measuring every relevant movement of the entire body, as well as provide tactile feedback to different parts of the body. The least advanced forms measure only the acceleration of one or very few parts of the body. Indeed, such limited forms of motion tracking are currently the only affordable VR technologies for most ordinary consumers. Off-the-shelf head-mounted displays, if they were to function optimally, would foster a presence in the virtual environment that is qualitatively different from the presence one experiences when the three-dimensional environment is projected onto a monitor and one changes perspective with the use of mouse, keyboard and/or gamepad.[17] However, the use of such head-mounted displays in virtual environments, and especially virtual *worlds,* has many flaws,

including time delay ('lag'), artificially narrow field of vision, reduced image resolution due to bandwidth limitations, and reports of considerable eye strain and nausea (Patterson, Winterbottom, & Pierce, 2006; Regan, 1995). If these flaws are sorted out, it is probable that this kind of augmentation of virtual worlds such as *Second Life* will become the most likely point of convergence between virtual worlds and virtual reality in the near future. The addition of multiple access to virtual reality, although still not feasible, clearly makes an important difference for the same reasons that virtual *worlds* are different from mere virtual *environments.* Having multiple access in VR would require body-like representations of all the participants, which would lead to simulation of even more lifelike activities, for instance as made possible by computer mediation of touch. However, since this technology is presently not feasible—at least not beyond a very limited degree—it hardly warrants a term of its own. Insofar as I will refer to such (future) technologies, I will simply use the qualifier *multi-access* virtual reality.

Definitions

After having reviewed some influential definitions of virtuality and analysed and re-phrased the distinguishing characteristics of the criteria found relevant—computer simulation, interactivity, indexicality, multiple access and first-person view—we get the following definitions:

> **Virtual x**: Interactive, computer-simulated x (or, x made possible by interactive computer simulation).[18]
>
> **Virtual environment**: Interactive, computer-simulated, *indexical* environment
>
> **Virtual world**: Interactive, computer-simulated, indexical, *multi-user* environment
>
> **Virtual reality**: Interactive, computer-simulated environment experienced from a *first-person view*.

The relation between these concepts can be further illustrated through a Venn diagram (see Figure 3-2).

In order to illustrate these definitions further, we can arrange a table of the relevant criteria as applied to some of the more popular forms of virtuality, including some non-virtual entities for comparison (Table 3-2).

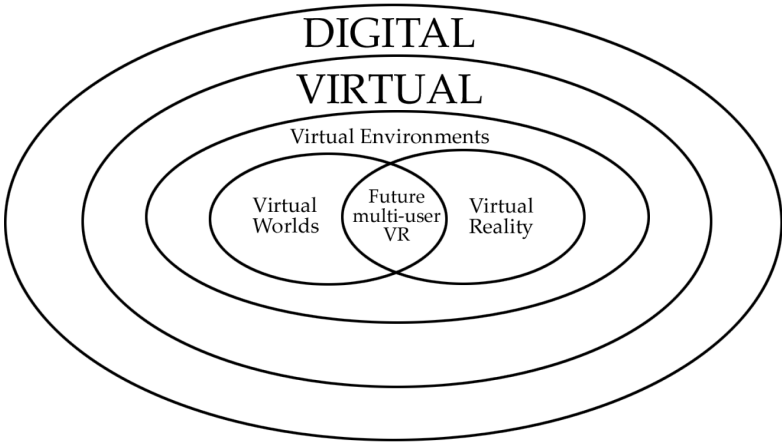

Figure 3–2: Illustration of the relation and intersections between the different categories

CATEGORY	COMPUTER SIMULATION	INTER-ACTIVE	INDEX-ICAL	MULTI-ACCESS	FIRST-PERSON VIEW	EXAMPLES
Multi-access VR	X	X	X	X	X	Future VR (including HDM-augmented virtual worlds when feasible)
Virtual Reality	X	X	X		X	CAVE, Virtusphere, Hi-tech flight simulators
Virtual World	X	X	X	X		*Second Life*, *World of Warcraft*, Sims Online
Virtual Environment	X	X	X			Most modern, off-line video games, incl. Nintendo Wii
Virtual Community	X	X		X		LambdaMoo, chat rooms, discussion forums, social networking sites
Virtual (generic)	X	X				Virtual banks, libraries, universities
Real-time streaming media	X					Digital TV channels
Stereoscopy	(X)		X	X	X	3D movies, holograms
Physical reality		X	X	X	X	
Dreams and illusions		X	X		X	

Table 3–2: Taxonomy of 'virtuality' and its subcategories

This table requires some further elaboration. The combination of first-person view and multiple access (to such an extent that it requires networked communication) is currently not feasible due to technological limitations. Thus, multi-access is not a necessary criterion for virtual reality, which thereby includes high-end technologies such as the CAVE and Flight simulators. As explained above, virtual worlds and virtual environments differ from virtual reality in that they do not require a first-person view, and virtual worlds differ from virtual environments in that they allow multiple access. Furthermore, virtual communities differ from virtual environments in that they do not require indexicality, but they are qualitatively different from virtuality in general, which does not require interaction between multiple users. Finally, the 'virtual' category, which encompasses all the other subcategories, requires computer simulation and interactivity, which *together* are necessary and sufficient criteria. That is, an entity is virtual if and only if it requires an interactive computer simulation for its existence. 'Required for its existence' corresponds to how for instance 'social entities' require a social collective for their existence.

I have also included a number of non-virtual entities for comparison. As mentioned, there is no qualitative difference between streaming media and traditional media, since none of them involve interactivity (at least not to any interesting degree). Three-dimensional movies and the like are three-dimensional, can be shared with multiple users and are seen from a first-person view, but they still lack interactivity (hence only limited indexicality) and consequently do not differ substantially from other forms of visual illusions. In order to illustrate how the virtual occupies a middle ground between dream and reality, as elaborated above, the table shows that reality is not a computer simulation,[19] and that dreams and hallucinations lack the consistency, coherence and intersubjective availability supplied by both physical laws and computer simulations, and the presence of multiple users. The important point is that dreams and hallucinations do not allow for 'multiple access', in contrast with both reality and virtual worlds.

It might seem odd that in the figure above, physical reality has *fewer* properties than a multi- access VR, but this merely entails that if all of physical reality were to be simulated in a virtual reality, the *addition of* computer simulation will make the virtual reality different from physical reality—i.e., the differences between the two would be a result of the fact that the interactivity, indexicality, multi-access and first-person views only exist as a result of the computer simulation. In other words, claiming that virtual reality *lacks* something is entirely consistent with the figure above—the lacks might be a result of the addition of the underpinning computer simulation.

Why are these distinctions important?

Although precision and clarity is important in general, it is particularly important when we make claims about the ethical and societal implications of technology—even more so when we make claims about 'virtuality', because such claims often involve a comparison between virtual and actual. This relates to the so-called *principle of formal equality*,[20] which states that a difference in treatment or value between two domains can only be justified on the basis of a relevant and significant difference between the two. For instance, many issues in animal ethics can be approached by first discussing the theoretical differences between humans and other animals (e.g., that some animals have a significantly less developed Central Nervous System) and subsequently discuss to what degree these differences are relevant to their value (e.g., that the less developed CNS indicates little or no ability to experience pain). More generally, for something to be of more or less value than something else, the two must be significantly different in one way or another, and we must be able to justify why this difference entails a difference in value. This entails that when we ask how trust in virtual worlds may differ from trust in actual reality, we are asking 1) how do virtual worlds differ from actual reality, and 2) how do these differences determine the possibility for trust?

A few examples may illustrate this point. Howard Rheingold argues that the lack of spontaneous bodily gestures and facial expression is the reason for the 'ontological untrustworthiness' of virtual acts of communication (Rheingold, 2000, p. 177). This critique seems to hold for virtual communities where there is no bodily representation, but it may not hold for virtual *reality* where motion tracking can relay some spontaneous bodily gestures and even facial expressions. Similarly, Lucas Introna argues that 'virtual communities are… different to those that are situated, embodied and collocated in that they have much less resources available to express and secure their identity through shared community objects' (Introna, 2007, p. 170). Again, this seems to hold for many types of virtual communities, but perhaps not for virtual *worlds*, since the multi-user requirement entails that these virtual worlds have a completely different level of permanence, which may in turn entail the possibility of having relatively stable virtual entities that *can* act as shared community objects. It is not my purpose to discuss the validity of these claims here, but to point out that the relevant considerations will necessarily be different depending on what type of virtuality we are talking about—and it has been the purpose of this paper to contribute a set of definitions and distinctions that can be useful in this regard.

Conclusion

The purpose of this paper has been to propose definitions for the terms 'virtual', 'virtual environment', 'virtual world', and 'virtual reality' as well as to clarify how they relate to each other. I have also argued in more detail how important it is to emphasize the role of the interactive computer simulation as provider of regularity, intersubjective availability (in the case of virtual communities and worlds), and hence an epistemological objectivity that in many ways is equivalent to that of the actual world. Although this may not be the final word on how these concepts ought to be defined, I hope to have shown the importance of providing such definitions, the complexities involved in doing so, as well given a glimpse into the fundamental philosophical challenges that are raised by the characteristics of different forms of virtuality. Clarity in these issues is of utmost importance when we move on to make claims about the epistemological and ontological status of virtual entities, such as when discussing the nature and possibility of 'trust' in virtual environments, virtual worlds or virtual reality.

Acknowledgments

I would like to thank Charles Ess for carefully reading and commenting on this paper, as well as Philip Brey and Adam Briggle for giving feedback on earlier versions. This paper has been presented at the workshop 'Philosophy of Virtuality: Deliberation, Trust, Offence and Virtues' (NTNU, 2009) and the 'Philosophy of Computer Games Conference' (UiO, 2009) and I am grateful for the helpful feedback from both audiences.

Notes

1. See Marianne Richter (this volume) for an overview and discussion of ways in which 'actuality', 'reality' and related terms have figured in the history of philosophy.
2. The term 'cyberspace' was coined by William Gibson in his famous *Neuromancer* (1984). The term 'infosphere' appears to have been coined by R.Z. Sheppard (Sheppard, 1971) and plays an important role in Luciano Floridi's philosophy and ethics of information (see e.g., Floridi, 2002).
3. Cf. Søraker (2010) for a more thorough review and critique of a number of other suggested definitions. I single out Stanovsky, Brey and Burdea & Coiffet here because the discussion of their definitions will be directly relevant for my proposed definitions later.
4. The CAVE is an immersive virtual reality where the images are projected on the walls of a cube surrounding the user. Such systems usually do not employ networked communication due to the immense speed and bandwidth that would be required, which also illustrates why

networked communication is not an apt criterion for 'virtual reality'. Cf. http://www.evl.uic.edu/pape/CAVE/ [Retrieved May 22, 2010].

5. The notion of congruence recurs throughout KdRV, but see in particular B278-279 ('congruence with the criteria of all actual experience'), A376, A112 and A451/B479. In the words of Kant, the 'connection of appearances determining one another with necessity according to universal laws [is] the criterion of empirical truth, whereby experience is distinguished from dreaming' (KdRV, A493/B521).
6. Cf. DoD Modeling and Simulation Glossary (United States Department of Defense, 1998, p. 124).
7. The importance of real-time simulation will become clearer in defining virtual worlds as essentially 'multiple access' below.
8. The term 'avatar' as referring to virtual agents was popularised by Neal Stephenson in his influential novel *Snow Crash* (1993). The term originally comes from Hindu and refers to the physical incarnation of a divine being.
9. As Espen Aarseth has reminded me, in some computer games you interact with the virtual environment from a God's-eye point of view and/or control a collection of 'avatars', e.g. in real-time strategy games such as Command & Conquer. Although these environments lack indexicality as defined above, they do embody three-dimensionality and many of the other features of virtual environments, thus I suggest referring to these simply as non-indexical virtual environments or, if massively multiplayer, as non-indexical virtual worlds.
10. There are, or rather were, non-graphical, one-player environments (e.g., text-based adventure games or 'interactive fiction', such as *Galatea*) that embody many of the same properties as virtual environments, including a weak form of indexicality. These are relatively rare, and increasingly so, hence do not require a term of their own and can be referred to as *non-graphical* virtual environment. If such non-graphical environments allow for multiple users (e.g., online text-based role-playing games such as LambdaMOO), they will be covered by the 'virtual community' category as explained below.
11. Another form of involving multiple users is sometimes referred to as 'hot seat'. This means that a number of participants can participate in the same environment without the use of networked communication. This is, for natural reasons, limited (usually no more than four participants at the same time) since the view will be the same for all participants. Hence, I will not include these technologies under the heading of virtual worlds.
12. I will return to the possibility of multiple users in fully immersive virtual reality below.
13. Indeed, 'world' is made up of the roots of *wer* and *ald*, originally from German, meaning 'man' and 'age'. Thus, the old English 'world' literally meant 'age of man'.
14. See the Association of Virtual Worlds' *Blue Book* (Association of Virtual Worlds, 2008) for a comprehensive list of the many hundred different virtual worlds currently in use.
15. In the distant future, it is not impossible that we will see brain-computer interfaces that feed the visual stimuli directly to the brain instead, as illustrated in the Matrix movies. Indeed, there have been some early attempts to control *Second Life* by means of a brain-computer interface (Keio University, 2008).
16. In continuation of the quote above, Husserl adds that 'the like holds, obviously, for touch' (ibid).
17. There are also other forms of motion-tracking that foster some degree of immersion, and as such constitutes a grey area between virtual environments and virtual reality. To take but one

example, Microsoft's upcoming release of 'Project Natal', an addition to the XBOX console, promises interaction solely through bodily gestures and speech recognition. These kinds of technologies allow for a concrete type relationship between your body and your avatar, since the avatar's action will typically correspond only to the type, not the token, activity of your body. Currently, this does not count as virtual reality proper, however, since it lacks a pure first-person *view* as described.

18. This definition may or may not exclude some entities commonly referred to as 'virtual'. It intentionally excludes things like 'virtual radio', which is a misnomer (see above). It also excludes terms like 'virtual memory', 'virtual server' and other computer science terms. These are established terms and I simply grant that they use the term 'virtual' in a different sense from the way I have defined it above. 'Virtual memory' and the like seem to rest on the common sense definition of 'almost the same as', or 'fulfilling the same role as' actual computer memory, which can be very misleading when talking about virtual worlds and entities therein, as discussed above.
19. Although it is logically possible that reality is computer-simulated, as argued by Nick Bostrom (2003), this is not an option I will entertain in this paper.
20. The principle is usually attributed to Aristotle (*Nicomachian Ethics,* V.3. 1131a10-b15; *Politics*, III.9.1280 a8-15, III. 12. 1282b18-23; cf. Gosepath (2008)). See also Søraker (2007) and Wetlesen (1999).

References

Aarseth, E. (2006). Doors and Perception: Fiction vs. Simulation in Games [Electronic Version]. *Digital Arts & Culture Conference*. Retrieved September 2, 2009 from http://www.luisfilipeteixeira.com/fileManager/file/fiction_Aarseth_jan2006.pdf.

Association of Virtual Worlds. (2008). *The Blue Book: A Consumer Guide to Virtual Worlds*. Retrieved October 25, 2008, from http://associationofvirtualworlds.com/thebluebook/the_blue_book_august_2008_edition.pdf.

Borgmann, A. (1999). *Holding On to Reality: The Nature of Information at the Turn of the Millennium*. Chicago, IL: University of Chicago Press.

Bostrom, N. (2003). Are You Living in a Computer Simulation? *Philosophical Quarterly, 53*(211), 243–255.

Brey, P. (1998). Space-Shaping Technologies and the Geographical Disembedding of Place. In A. Light & J. M. Smith (Eds.), *Philosophy and Geography III: Philosophies of Place* (pp. 239–263). Lanham, MD: Rowman & Littlefield.

Brey, P. (1999). The Ethics of Representation and Action in Virtual Reality. *Ethics and Information Technology, 1*(1), 5–14.

Brey, P. (2003). The Social Ontology of Virtual Environments. *American Journal of Economics and Sociology, 62*(1), 269–282.

Brey, P. (2008). Virtual Reality and Computer Simulation. In K. E. Himma & H. T. Tavani (Eds.), *The Handbook of Information and Computer Ethics* (pp. 361–384). Hoboken, NJ: John Wiley & Sons.

Floridi, L. (2002). On the Intrinsic Value of Information Objects and the Infosphere. *Ethics and Information Technology, 4*(4), 287–304.

Gibson, W. (1984). *Neuromancer*. New York: Ace Book.

Gosepath, S. (2008). Equality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*.

Hume, D. (1739/2008). *A Treatise of Human Nature*. Sioux Falls, SD: NuVision.

Husserl, E. (1989). *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy [Second book]* (R. Rojcewicz & A. Schuwer, Trans.). Dordrecht, Netherlands: Kluwer.

Introna, L. (2007). Reconsidering Community and the Stranger in the Age of Virtuality. *Society and Business Review, 2*(2), 166–178.

Kant, I. (1781/1997). *Critique of Pure Reason* (P. Guyer & A. W. Wood, Trans.). Cambridge: Cambridge University Press.

Keio University. (2008). Press Release: Keio University Succeeds in the World's First Demonstration Experiment with the Help of a Disabled Person to Use Brainwave to Chat and Stroll through the Virtual World [Electronic Version]. Retrieved August 1, 2008 from http://www.keio.ac.jp/english/press_release/080605e.pdf.

Lanier, J. (1999). Virtual Reality. *Whole Earth* Fall 1999. Retrieved January 6, 2008, from http://findarticles.com/p/articles/mi_m0GER/is_1999_Fall/ai_56457593

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Basil Blackwell.

Lopes, D.M. (2001). The Ontology of Interactive Art. *Journal of Aesthetic Education, 35*(4), 65–81.

Lopes, D.M. (2004). Digital Art. In L. Floridi (Ed.), *Blackwell Guide to the Philosophy of Computing and Information*. Oxford: Blackwell.

Malebranche, N. (1688/1997). *Dialogues on Metaphysics and on Religion* (N. Jolley & D. Scott, Trans.). Cambridge: Cambridge University Press.

North, M.M., North, S.M., & Coble, J.R. (1998). Virtual Reality Therapy: An Effective Treatment for Phobias. *Stud Health Technol Inform*(58), 112–119.

Patterson, R., Winterbottom, M.D., & Pierce, B.J. (2006). Perceptual Issues in the Use of Head-Mounted Visual Displays. *Human Factors: The Journal of the Human Factors and Ergonomics Society, 48*, 555–573.

Regan, C. (1995). An Investigation into Nausea and Other Side-Effects of Head-Coupled Immersive Virtual Reality. *Virtual Reality, 1*(1), 17–31.

Rheingold, H. (2000). *The Virtual Community: Homesteading on the Electronic Frontier*. Cambridge, MA: MIT Press.

Robbins-Bell, S. (2008). *Using a Faceted Classification Scheme to Predict the Future of Virtual Worlds*. Paper presented at the Internet Research 9.0 (AoIR) conference, October 15–18, 2008. Copenhagen, Denmark.

Rosen, G. (2006). Abstract Objects. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Spring 2006 Edition)*.

Searle, J. (1995). *The construction of social reality*. New York: Free Press.

Shannon, C.E., & Weaver, W. (1949/1998). *The Mathematical Theory of Communication*. Urbana: University of Illinois Press.

Sheppard, R.Z. (1971). Rock Candy [Electronic Version]. *Time*. Retrieved July 12, 2009 from http://www.time.com/time/magazine/article/0,9171,905004,00.html.

Sherman, W.R., & Craig, A.B. (2003). *Understanding Virtual Reality*. San Francisco: Elsevier.

Shields, R. (2003). *The Virtual*. London: Routledge.

Søraker, J.H. (2007). The Moral Status of Information and Information Technologies—a Relational Theory of Moral Status. In S. Hongladarom & C. Ess (Eds.), *Information Technology Ethics: Cultural Perspectives* (pp. 1–19). Hershey, PA: Idea Group Publishing.

Søraker, J.H. (2010). *The Value of Virtual Worlds and Entities [Dissertation]*. Enschede: University of Twente.

Stanovsky, D. (2004). Virtual Reality. In L. Floridi (Ed.), *The Blackwell Guide to the Philosophy of Computing and Information* (pp. 167–177). London: Blackwell.

Stephenson, N. (1993). *Snow Crash*. New York: Bantam Books.

Steuer, J. (1992). Defining Virtual Reality: Dimensions Determining Telepresence. *Communication in the Age of Virtual Reality, 42*(4), 73–93.

Tugendhat, E. (1976). *Vorlesungen zur Einführung in die sprachanalytische Philosophie*. Frankfurt Am Main: Suhrkamp.

United States Department of Defense. (1998). DoD Modeling and Simulation (M&S) Glossary, DOD 5000.59-M. Retrieved March 10, 2007, from http://www.dtic.mil/whs/directives/corres/pdf/500059m.pdf

Wetlesen, J. (1999). The Moral Status of Beings Who Are Not Persons: A Casuistic Argument. *Environmental Values, 8*, 287–323.

Wetzel, L. (2007). Types and Tokens. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Winter 2007 Edition)*.

Wyller, T. (2000). *Objektivitet og jeg-bevissthet: en aktualisering av Immanuel Kants filosofi*. Oslo: Cappelen.