

The transhumanism project will gain momentum with advances in technology, in basic science and in philosophy, as well as in bioethics. However, there are minefields that jeopardize this progress – one such minefield is a fundamental problem in pure philosophy: fictional entities and how we refer to the nonexistent. In the absence of solutions to the problems that arise in this area of philosophy, progress in the technology necessary for augmented reality will be considerably impeded. I will argue there are forms of augmented reality that are metaphysically impossible and that believing that such forms are possible (both metaphysically and physically) creates a form of skepticism.

Introduction: augmented reality in the near future

Imagine the following scenario, which you might think may well happen in the not too distant future. Fred is under orders from his doctor not to indulge in foods that will elevate his already dangerously high cholesterol level. Unfortunately, Fred has little will power to resist the delicious foods he is forbidden to eat and that may well kill him if he does eat them. Luckily for Fred, it is now possible to wirelessly connect his cerebral cortex directly to the Internet. He connects to a website that caters to weak-willed, high cholesterol people. Instantly a tiny elfin figure –

Peter Pan – appears in Fred's visual field. Peter Pan engages in repartee with Fred, after which he then saunters over to the overstuffed pastrami sandwich Fred is desperately trying to ignore. He pulls down his drawers and defecates on the sandwich. Fred recoils in disgust, murmurs to himself that Peter Pan is standing on his sandwich, but that Peter Pan does not exist, and tosses

the sandwich into his waste disposal unit. Peter Pan disappears and Fred settles down to eat his AMA-approved lunch of lettuce and cucumber.

This scenario poses a number of difficult philosophical puzzles about what actually happens in it – is Fred having a conversation with a fictional entity? Is a fictional entity walking on top of Fred’s sandwich? Can Fred truly say both that Peter Pan is standing on his sandwich and that Peter Pan does not exist? (If we accept Saul Kripke’s views on fictional entities, Fred cannot truly say that.) In augmenting Fred’s visual field with beings that are fictional, but which are visible to him, what exactly is being done? Augmented reality is a vision for the near future that will require, if it is to happen, a detailed understanding of brain processes and brain-machine interfaces. It might take the form of direct wireless connection of the human brain to the Internet or virtual reality devices, such as virtual reality glasses. There are multiple important functions subserved by such technology. For one, the uses of virtual beings that appear in one’s visual field are almost endless. A virtual companion could serve as a guide when traveling in a new city, or could help one in resisting the attractions of behaviors that are either dangerous or that will create health problems, or could be a dedicated tutor or a narrator of community events or a story-teller or an entertainer. They could be extensions of one’s bodily appendages, such as hands, nose or feet – a surgeon might use one to help guide her fingers during an operation or to mark places in bodily tissues that must be avoided or a golfer might use one to help move the club in a desired trajectory.

One might think that the only problems for implementing augmented reality are issues concerning neurophysiology and brain-machine interfaces. That these are problems for augmented reality is certainly true, but it is not the complete story. There are additional problems that may call into question the very idea of an augmented reality. These come from an area of pure philosophy, namely reference to fictional entities and the logic of nonexistence. I will argue that if certain views about fictional entities and the logic of nonexistence are accepted – those of Saul Kripke – then certain kinds of augmented reality are metaphysically impossible. If so, it follows that we might be wrong about what we think we are observing when virtual fictional beings appear in our visual field or when there is any kind of communication between fictional virtual beings and real beings. So we will need to distinguish between different kinds of reality augmentations, some of which are philosophically incoherent and some of which are not.

Kinds and degrees of reality augmentation

The term “augmented reality” marks off a large parcel of possible virtual augmentations. Although it is not the point of this paper to provide an analysis of the concept of an augmented reality, nor to describe all the different kinds and respective degrees of augmented reality, it is worthwhile to briefly sketch some of the kinds. The claim of this paper is that one kind of augmented reality is philosophically incoherent. That is, there are a priori reasons to believe that it cannot happen. It is not that the concept makes sense, but is either physically or technologically unachievable. Rather, the very concept is incoherent. It is metaphysically impossible. Thus, there is nothing that could be physically or technologically implemented. The importance of the claim is that it shows how we could be radically deceived about what we think we have engineered. That is, certain reality augmentations involving fictional virtual beings look like they are technologically achievable, but that is the mistake. It is not those reality augmentations, but reality augmentations of a different kind that appear to be reality augmentations involving fictional beings that are technologically achievable. To repeat: I am not arguing that a particular kind of augmented reality is either physically impossible or technologically impossible (though physically possible). My claim is not about limitations of the physical world, but rather about the

concept of reality augmentation and the metaphysical limitations imposed by a philosophical theory of fictional entities.

Hearing aids are a simple and well-known kind of reality augmentation. Incoming sound waves are amplified for people whose hearing mechanism has been compromised in some way. Similarly, eyeglasses are a form of reality augmentation. In both cases, incoming signals are technologically altered so that an underlying physical impairment in the sensory mechanisms sustaining a cognitive ability (such as sight) is overcome. Adding machines, calculators, slide rules and computers (for making numerical computations) are ways of augmenting our arithmetical skills, but are not reality augmentations, since it is not some feature of the physical world (our reality) which is being augmented. Of course, one might argue that if one takes augmentation of reality to be either adding something to, changing something in, or subtracting something from reality, then enhancing one's ability to add numbers by use of a calculator might be seen as a form of reality augmentation, since one is adding to "reality" sums which one would not have produced without the use of the calculator. I will not worry whether or not this way of bringing calculators (and much else) under the umbrella notion of augmented reality is conceptually appropriate, since nothing in what follows hangs on a resolution of this issue.

The term "augmentation" suggests an enhancement, though the dictionary counsels neutrality on that: to augment is to become greater in size, extent or quantity. Certainly, some reality augmentations will be seen as enhancements (headphones) while others will be seen as treatment of some underlying pathology (hearing aids). Since this distinction is rather difficult to draw and may be hostage to our definition of normal health,¹ it is best not to appeal to it in elaborating the concept of augmented reality. We do well, though, to focus on the notion of "reality." If reality is defined to be the totality of all that is real, then it is not clear how reality can be extended. If reality includes all and only what is real, then any extension of it must consist of things that are not real, and thus not part of reality. In which case, there is, by definition, no such thing as augmented reality. One way out of this definitional impasse is to define reality to be a complex psychological state that involves human perception and human experience. The methodological point, then, is that what is an augmentation of reality critically depends upon the conception of reality to which we appeal.

One way to think of reality is as the sensory and cognitive world that we inhabit. On this view, reality is not some transcendental realm forever hidden from us that lies below our sensory and cognitive apparati, but is just what we experience in virtue of the activities of our sensory and cognitive apparati. Reality is, on this view, what we experience, and a description of the phenomenology of our experiences would be a description of the fabric of reality (though not a description of the physical structure and physical laws of reality). Adding to or modifying what we experience, then, is augmenting our experiences and thus augmenting reality. Since our experiences can be either veridical, on the one hand, or hallucinations or illusions, on the other, one might insist that reality is just the set of our veridical experiences, since hallucinations and illusions are experiences of what is not real. But this move threatens to restore the definitional impasse it was the point of defining reality in terms of experiences to remove. Instead, I will stipulate that augmented reality consists of any change in the totality of our sensory and cognitive experiences (at any given moment) that is produced by some form of technology. The three basic kinds of technologically mediated changes are additions, modifications and deletions. Another stipulation is the following: how the world appears to be as it is experienced by a human being is called one's total experiential field. This is a term borrowed from Gestalt psychology (Gurwitsch 1964, Chapter 1).

One kind of augmented reality, on this view of reality, is a virtual modification in which various features of reality are virtually enhanced. For instance, putting on rose-colored glasses produces an augmented reality in which the natural colors of objects are modified – and perhaps enhanced. Listening to music piped into one’s ears by a headset, while walking, is a form of augmented reality in which features are added to reality. One adds to the natural sounds of the local environment the sounds of, say, a Beethoven piano sonata. In such cases, natural sounds in the local environment might be masked, modified or deleted. Olfactory and tactile augmentations of reality are also technologically available. If we stipulate that a rose-colored world is non-normal, then we can distinguish between normal and non-normal augmentations of reality. We can also distinguish between what exists in reality and what does not. Reality augmented with an image in one’s visual field of a non-existent flower with natural colors is a normal addition to it of something that is non-existent. Of course, the image of a non-existent object exists in one’s visual field and is, in that sense, a real part of one’s visual field.

Given a sensory component of the human brain, there are reality augmentations of it. What, however, of virtual linguistic inscriptions, such as words, that augment one’s visual field? (The importance of this kind of reality augmentation, and its associated problems, will be developed below.) These inscriptions are not images of what exists in the real world (unless they are, either intentionally or coincidentally, images of inscriptions that exist in the real world and which occupy the same place in one’s visual field that the image of the real world inscription would occupy – though it may seem bizarre to replace an image of an object with another image that is indiscernible from it). This kind of augmented reality is not now available, but in the near future it will be technologically possible to mark objects in one’s visual field with linguistic inscriptions. Similarly, visual objects could be virtually marked with olfactory, aural, tactile, or gustatory signals. Such multi-modality sensory markings would be inscriptions, though, only if they provided information of some sort (that it is the function of an inscription to provide). Providing useful information to the user is certainly a primary function of future augmented reality technologies. Although it is important to distinguish between the manifold ways of representing information in a user’s total experiential field, devising such a catalog will not be our concern here. Suppose virtual non-existent beings are added to one’s visual field. If they talk, move, touch, and so on, those features will be features of the total experiential field that constitutes one’s reality. More importantly, if they provide information to a user by speaking, the user will be able to employ that information in performing various cognitive activities – or so our intuitions tell us.

Augmented reality and the problem of experiential field integration

In a future in which it is technologically possible to connect the human brain directly to the Internet, one form of augmented reality would be one in which items of various sensory and cognitive modalities from the Internet are integrated into one’s total experiential field. Unlike the cinema (though it certainly could include cinema-like augmentations), where the images projected onto the movie screen occupy a rectangular surface in the center of one’s field of vision, images and linguistic symbols would be scattered throughout one’s field of vision. Some of the images will be of virtual objects. One problem with an augmented reality of virtual objects is how they can interact in a natural way with the natural objects (including beings), natural features, and natural properties that are represented in one’s total experiential field. For instance, suppose that a virtual being that looks like a real person is in Smith’s total experiential field. This virtual companion could be helpful when Smith visits a new city. The companion would provide valuable information to him about city life. One can even foresee virtual beings serving as companions for people visiting new places, where it would be dangerous to seek temporary companionship. It might, however, be dangerous to Smith for a virtual figure to occupy some

portion of his visual space. For instance, if a virtual figure occupying Smith's visual space is analogous to simply pasting a picture of a man on Smith's glasses, then crossing a street in heavy traffic will be an invitation to disaster. How can virtual beings be integrated into one's visual field in such a way that valuable information about the local environment is not lost? That is a hard problem, and solving it will require an exhaustive knowledge of the human mind/brain at both computational and neurophysiological levels.

There are integration problems in other sensory modalities. If the virtual object is a virtual being who talks, the sounds it makes must not destroy or mask the natural sounds from one's local environment. Here, too, sounds provide valuable information necessary for survival. Losing sounds is also a recipe for disaster. We will assume that the problem of integrating perceptual signals into one's total experiential field has been solved. The world is a complex and continually changing place where what is true at one time may no longer be true a half-second later. Integration of virtual fictional or non-fictional beings into one's visual field would require, minimally, a solution to the frame problem in artificial intelligence. There are some philosophers who have argued that there will never be a computational solution to the frame problem. Whether that is or is not the case, the claim of this paper is not affected. These claims are not hostage to the technological means by which augmented reality is achieved.

There are certainly problems of information overload that need to be addressed (besides the problems of sound integration) in augmented reality technologies. It would be a worthless technological project if Smith acquired so much information from the combination of his local environment and reality augmentations that he could not store all of it in short-term memory. (Decisions would then have to be made about which information should be stored and which not, but that is a problem that does not affect the claim that is argued in this paper.)

There is, though, a different integration problem for sounds. Suppose that the virtual being speaks to Smith. There are situations in which it would be important that it appears to Smith, phenomenologically, that the virtual being is speaking to him and not simply speaking to a general audience (in the way in which a carnival barker speaks to anyone and everyone). It might be thought that this problem is a technological one. However, that is not so, if certain circumstances obtain. If the virtual being augmenting one's total experiential field is a fictional being, it would not be possible for it to speak to Smith. It might appear to Smith that this is precisely what it is doing, but Smith would have a false belief if he believed that to be the case. In order to see just what the problem is, we need to consider two different kinds of virtual beings inhabiting one's total experiential field. The first kind is a fictional being, while the second kind is non-fictional: it is a conduit for information that is transmitted from another human being or is information that is transmitted in virtue of a program being executed (and where there may or may not be program designer intentions in place). Let's consider the second kind of virtual being first. Just as we acquire information from other people from (but not restricted to) their words (say, in a conversation), we acquire information from virtual non-fictional beings in the same way. There is a causal connection of us to the virtual non-fictional being and of the virtual non-fictional being to some source of information. Given that the causal connections are in place, if the virtual non-fictional being tells Smith that Jones wants to meet with him at the Hilton, then the referent of the word "Jones" is the woman named "Jones," and the referent of the word "Hilton" is a certain hotel that is part of the Hilton chain of hotels. (The same is true if the information is transmitted in virtue of a program being executed, though the causal chains might be more complex in this case.)

If the virtual non-fictional being augmenting Smith's visual field is causally connected to a real human being (say, Jackson) and it is not Jackson's intention that the virtual being be a fictional

character, then what Smith hears from the virtual non-fictional being is no different, in terms of its information content, from what Smith would hear from Jackson. (I ignore such things as tone of voice, emphasis, facial expressions, and so on. All of these can provide information in addition to the actual words employed in a conversation.) That is, Smith acquires the information that Jones wants to meet him at the Hilton whether it is communicated to him by the virtual being or by Jackson.

Augmented reality with virtual fictional beings

The first kind of virtual being – a virtual fictional being – creates special problems for communication different in kind from any of the problems created by virtual non-fictional beings. To fix ideas, suppose that it is possible to download a fiction (perhaps a film) into one's total experiential field in such a way that all of the problems discussed above have been resolved. There is another problem, and it has to do with the problem of reference to fictional entities and the logic of nonexistence. To put it crisply, there is an impenetrable barrier between real beings and fictional beings. The barrier is not physical. It is metaphysical, in the following way: it is a metaphysical impossibility for beings in reality to communicate with beings in a fiction and for beings in a fiction to communicate with beings in reality. That much might strike anyone as being trivially true. What we need to see is how it falls out of one philosophical program of how reference to fictional entities works. My argument in this paper is relativized: I will help myself freely to the claims made in only one philosophical theory of reference to fictional entities – Saul Kripke's theory (advanced in, among other places, Kripke 2011). It is not my aim in this paper to critically examine Kripke's views, to investigate the metaphysics of fictional entities, or to consider the problem of virtual fictional beings in augmented reality through the lens of different philosophical theories of reference to fictional entities.

The primary claim I argue for in this paper is that reality augmentation with fictional virtual beings is metaphysically impossible. It is impossible in two distinct ways. First, it is metaphysically impossible to communicate from real beings to fictional beings and, conversely, from fictional beings to real beings. Second, it is metaphysically impossible to integrate into a total experiential field both real objects and fictional objects (where the objects of either kind include, of course, beings). To see why both of these metaphysical impossibilities hold we need to examine Kripke's views on reference to fictional entities.

The problem of reference to fictional entities and to non-existent objects was thought to deliver a fatal blow to Millian theories of reference and to vindicate Russellian theories of reference. Kripke had destroyed Russell's theory of descriptions as an account of reference in 1970 (in his lectures, *Naming and Necessity*, at Princeton (Kripke 1982)). He devised an ingenious account of reference to fictional entities in (Kripke forthcoming 2012, some of which is anticipated in Kripke 1982). The central idea in his account is that in works of fiction it is only pretended that the conditions described actually obtain. That is, "a work of fiction is a pretense that what is happening is really going on" (Kripke 2011, 52-74). Names used within a fiction do not have referents. Rather, it is merely a pretense that they have referents. Thus, an author of a fiction can pretend that any theory of reference is in place and is satisfied within their fiction. An author can pretend that Russell's theory of descriptions is an entirely satisfactory account of the reference relation and pretend that reference is secured in her fiction by the appropriate descriptions. That this is a radical proposal goes without saying, since it implies that all of the propositions in a piece of fiction are not genuine propositions. Rather, they are pretend propositions. Propositions within fiction are not genuine because, if they were, we would be able to speak of counterfactual situations in which they are true. Suppose that there is a proposition about a particular fictional person in a work of fiction. It makes no sense to ask whether that proposition would be true under

various counterfactual conditions, since the fictional entities in the fiction have no essential features that could be used to determine whether they are, indeed, objects in the counterfactual situation. What criterion would we be able to use to pick out a fictional character in a given counterfactual situation?

Already in *Naming and Necessity*, Kripke discusses unicorns and explains why it is that these mythical creatures could not be discovered to exist (Kripke 1982, 24, 157-158). There are component epistemological and metaphysical theses that are somewhat connected. The epistemological thesis is that we would not be able to discover unicorns existed, because the mythical unicorn has no microstructure (such as DNA) which could be used to identify an actual creature as an instance of a unicorn. The metaphysical thesis is that there are no possible circumstances under which unicorns would have existed. The epistemological thesis depends on the metaphysical thesis: without an essence, there are no actual beings satisfying the description of a unicorn. Thus, without an essence, there is no metaphysics of the unicorn, and so there is no criterion by which we could epistemically identify a being as a unicorn. Further, there is no criterion by which we could say of a being in a counterfactual situation that it is a unicorn.

Consider the fictional being Hamlet. Shakespeare is merely pretending that the name “Hamlet” refers to a certain man and is pretending that he possesses features, such as certain psychological traits. There are many sentences in the play *Hamlet* that are about the person named “Hamlet.” But none of these expresses a genuine proposition, since it is metaphysically impossible to say under just what counterfactual circumstances such propositions would be true. In which case, it follows that we cannot say under what circumstances Hamlet would have existed. If so, it is not the case that fictional entities are possible, but non-actual, objects (since to be a possible object we must be able to say under what circumstances it would have existed). What if we want to say something about what happens in the play *Hamlet*? Do we simply say that the ghost of his father appears at the beginning of the play? Or is it that we can’t say any such thing, because we cannot express a genuine proposition about the appearance of the ghost of Hamlet’s father? If we say that Hamlet saw the ghost of his father, we must contextualize that assertion to the fiction (or employ an operator “in the fiction”). That is, we say that in the fiction it is true that Hamlet saw the ghost of his father. In the fiction, Hamlet is an actual and not a fictional person. What if we want to talk about the claims literary scholars have made about Hamlet? Kripke’s view is that when we talk about fictional entities from outside the context of the fiction in which they appear, fictional entities are abstract entities which are not mere ghostly entities. They are full-bodied abstract objects, in the same way that numbers (if such there be) are full-bodied abstract objects (Kripke forthcoming 2012, 18 Lecture 3). We can say, from outside the play, that the great scholar of Shakespeare, Gwynne Evans, considered *Hamlet* to be one of Shakespeare’s finest works. Moreover, it is true that this is so. Abstract objects can play roles in propositions that contribute to the truth-conditions of that proposition.

Kripke makes an analogy between fictional and real objects, on the one hand, and real and toy ducks on the other:

One thing which I think has caused confusion is the use of the word “real” here, where I contrast a fictional person to a “real” person. If the fictional person isn’t “real,” then after all there must be a weaker kind of existence involved. I don’t wish to make that kind of use of the word “real,” or that kind of comparison, when I speak of fictional characters. When I do this, a more appropriate comparison would be with the contrast between “real duck” and “toy duck.” Of course, there are differences as well as similarities. But just as a toy duck isn’t a real duck, though of course that doesn’t mean that the toy duck doesn’t really exist, so I want to say that a fictional person isn’t a “real” person, though that isn’t

to say that, in and of himself, he doesn't really exist, or isn't real, in the sense of "doesn't exist." On the contrary, just as there can be real toy ducks ... so Hamlet is a real fictional character ... (Kripke forthcoming 2012, 3 Lecture 2)

Fictional entities do not have a shadowy kind of existence. However, ordinary language is ambiguous with respect to the divide between real and fictional, and that creates problems, since we do not ordinarily preface our remarks with a description of the context to which they apply. In speaking of fictional entities it is easy to conflate one context with another. When we speak of what happens in a fiction we implicitly employ an operator "in the fiction" that we prefix to the sentence. Hamlet does not exist in the real world, but exists in a fiction. However, the fictional character Hamlet exists in the real world. It is false in the play *Hamlet* that Hamlet is a fictional character. But it is true in the real world that Hamlet is a fictional character. Notice that the truth-conditions of the sentence "Hamlet is a fictional character" are different in different contexts. It is false in one context (within the fiction), while it is true in another context (outside the fiction).

An objection

There is an objection to the preceding account of fictional entities and augmented reality. It is that actual agents *can* communicate with and learn from fictional entities. For instance, you can learn the Pythagorean Theorem from a fictional comic book character, Donald Duck. For another, if Donald Duck is in your augmented reality and he appears to talk to you, he *is* talking to you. How could he not be talking to you?

Although this objection has intuitive force, it fails. In the first place, fictional characters can make statements that are true, and one can learn from these statements. If we know it is the intention of the author of the fiction that her fictional characters utter actual world truths, then we are rational to take their utterances to be true. On the other hand, if we do not know the intentions of the author of the fiction, and do not antecedently know that the utterances of fictional characters in the fiction are true, we must independently confirm that they are true. In the second place, although it appears that a fictional character is talking to you, unless you know the intentions of the author of the fiction, you cannot assume it is talking to you. If the intention of the author of the fiction is that the fictional characters in it talk to a particular actual world person, then they are talking to the actual world person. One might even have a fiction in which the author allows for an actual world person to make responses to which a fictional entity in the fiction responds. Or one might have a fiction where the fictional characters address the reader of the fiction. But without knowing the intentions of the author of the fiction, one cannot assume that a fictional entity is talking to you, where it inhabits your augmented reality. There is certainly no causal connection between you and the fictional entity, although there might be a causal connection between you and a representation of that fictional entity. Similarly, there is no causal connection between you and the natural numbers, although there is a causal connection between you and a representation of the natural numbers, viz., the numerals.

However, one might also object that even if there is no communication between a fictional character and an actual person, it does not matter if an actual person takes there to be communication between them and fictional characters. I will argue in several sections below that doing this can create serious problems which might jeopardize the life of the user of an augmented reality. By the way, it might happen that an augmented reality technology allows a user to create her own fictions within her head. In which case, the user of the augmented reality is pretending that they are talking to the fictions in their augmented reality. However, since the user is the author, her intentions are transparent to her, and although her relations with her fictions are based on pretense, she can arrange things so that the fictions speak truths or speak to her.

Fictional entities in augmented reality: the cross-barrier communication problem

Jaakko Hintikka has argued that Descartes' cogito ergo sum argument is fallacious, since there are instances of it in which the premise is true, but the conclusion is false (Hintikka 1963). Hintikka chose as a counterexample the following: It is true that Hamlet thinks, but false that Hamlet exists. In which case, it is fallacious to infer that Hamlet exists from the premise that Hamlet thinks. In his Locke lectures, Kripke points out that Hintikka is incorrect, for he confuses the context in which we make a true statement about what happens in the play *Hamlet* with the context in which we, from the outside, talk about the play *Hamlet* (Kripke forthcoming 2012, 1-2 Lecture 3). From inside the fiction, it is both true that Hamlet thinks and that he exists. Outside the play, it is false that Hamlet exists and false that Hamlet thinks. Outside the play, it is true that the fictional character Hamlet thinks, but it is equally true that the fictional character Hamlet exists. Unless one's claims about fictional entities are subscripted with the context in which they are made, it is easy to make false claims about them.

If Shakespeare's fictional character Hamlet appears as a virtual fictional being in one's total experiential field, even if it appears to the user, phenomenologically, that Hamlet is speaking directly to him, it cannot be that this is really happening. If Hamlet says that he wishes anyone standing near him to leave the room, he is not talking to the user (even though the user appears to be standing next to Hamlet and even if there is no other being standing near Hamlet). There is a barrier to communication between the user and a virtual fictional entity, even though when virtual fictional entities appear in one's total experiential field, there is no physical barrier between them and the user of augmented reality technology. (In reading fiction, seeing a fiction in a film, or in a stage performance, we know what the barrier between real and fictional consists in.) It is a metaphysical impossibility for a virtual fictional being to communicate with the user of an augmented reality technology, even if, it appears, to the user, that communication is exactly what is occurring to the user. Suppose that the virtual fictional being uses the name of the user and says personal things about the user that only the user knows. No communication has taken place, since the virtual fictional being is not talking to the user.

Fictional beings in augmented reality: the foolish inference problem

Suppose that there is a real geographical location that a fictional character is described by its author as inhabiting. Let's say that the fictional character Hamlet lived in Denmark. If we do not pay attention to the context in which a statement about a fictional character is made, we can easily commit Hintikka's mistake. For instance, we might say Hamlet lived in Denmark and Hamlet appears in a play of Shakespeare, from which it follows (by existential generalization) that someone who lives in Denmark appears in a work of fiction (Kripke 1974, 224). It is true that in the fiction, Hamlet lives in Denmark. But unless the context is specified, it leads to the silly conclusion that there is someone living in Denmark who appears in a fiction!

Foolish inferences can be combated by not conflating the contexts in which they are made. In an augmented reality, though, I conjecture that it might easily appear to the user that there is one uniform reality and that the virtual fictional beings in it are just as real as the actual beings in it.² (Here there is the psychological problem of whether someone who could not distinguish virtual fictional beings from real beings would be said to suffer from some form of mental illness, such as having schizophrenic hallucinations.)

Consider an augmented reality in which there are multiple virtual beings, some of which are fictional, that function to help a human who has a motor impairment walk safely in a crowded

city. At some point all of the beings will appear as real to the user as the actual beings in the real world. Just as it is easy to become adjusted to inverting glasses, it might be easy to become “adjusted” to virtual beings in one’s reality enhancement. The user may easily become confused about the ontological status of the virtual fictional beings. He might come to believe that they are in the actual world and navigate accordingly. If so, he may well injure himself as a result of this false belief. Are there any fallacious inferences (based on a conflation of contexts) he might make that would jeopardize his safety? Even if he does not believe that the virtual beings (fictional or otherwise) are a part of the actual world, he might make many different kinds of fallacious inferences. For instance, he might infer that a virtual fictional being has real money, since it has a virtual wallet. Or he might infer that since the virtual being is really a fictional being, he really exists, in just the same way that his friend Harold exists. Or he might infer that the virtual beings have causal relations with all of the other objects in his total experiential field, in which case he must be careful that his actions do not have inappropriate causal consequences. There are many silly inferences that one can make; the point is that they can easily be made if contexts are ignored or conflated and there is no way in which the context conflation can be checked. Suppose that the deliverances of one’s inferences are given to cognitive or motor processes that are below the level of consciousness. If so, there is ample room for actions performed on the basis of the fallacious inferences that could jeopardize one’s life in multiple ways.

Fictional entities in augmented reality: the problem of authentic identification

Since a fictional being does not have a microstructure, there are no essential properties of it that can be used in determining under what conditions it would have existed. It follows that propositions describing fictional beings are merely pretend propositions. This creates a problem for the integration of virtual properties and real properties in reality augmentations involving fictions quite different from the integration problems we have already discussed. Suppose that an image of Hamlet is now in your total experiential field, walking in unison with you as you walk to the grocery store. You look at Hamlet as you walk and think that Hamlet looks rather worried and tired. If you think that you are thinking about Hamlet, though, you are dead wrong. You are not – even though it appears to you that this is exactly what you are doing. You are thinking about some image that bears a strong resemblance to what you think is the way the fictional character Hamlet looks.³ But since there are no identity conditions for the fictional character Hamlet, there is no way of determining whether the Hamlet that appears in your total experiential field is really Hamlet (unless it is the intention of the designer of the reality augmentation that it is the fictional character Hamlet). Similarly, if you encounter an animal that appears visually identical to a unicorn, you can’t say that it is a unicorn, since there are no identity conditions for unicorns. Only if you were, *per impossible*, to be in the fiction (of either the play *Hamlet* or the myth about unicorns) would you be able to tell if it was Hamlet or was a unicorn. But that cannot happen – even if you are the author of *Hamlet* or the author of the unicorn myth (unless it is the actual play *Hamlet* or the actual unicorn myth). It is not that you are in the play *Hamlet* when you are its author, but rather that you create the pretense and so are in a position to say what is and is not part of the pretense. However, if Shakespeare were transported to the twenty-first century and given a reality augmentation in which Hamlet is a fictional being appearing in his visual field – but it is not the fiction Shakespeare created (his play *Hamlet*) – then Shakespeare would not be able to say whether this fictional being appearing in his visual field is Hamlet or not (unless he knows the intentions of the designer of the reality augmentation).

This particular integration problem for fictional entities appearing in augmented reality might strike one as having no unwelcome consequences for any of the manifold uses of reality augmentations. That is, the user is simply wrong about taking the virtual fictional being to be Hamlet and “so what?” That is not true, though. One possible use of fictions in reality

augmentations would be to have the user experience what it is like to be in a certain, say morally troubling, situation described in the fiction. One of the ways in which fiction is important is in describing how characters extricate themselves from difficult moral problems – from which one can learn a great deal. One develops a moral sensitivity by reading good fiction, such as the novels of Iris Murdoch. Suppose that future technology has the means to project and integrate fictions from Iris Murdoch's novels into one's total experiential field. The hope would be that the user would be able to experience what it is like for the characters in her novels to engage in various forms of human experience by interacting with them on an individual basis. You interact with, say, Bradley Pearson, the protagonist of Murdoch's novel, *The Black Prince*. (This might also be used as a form of therapy for psychopaths or education in moral thinking for school children or for business executives.)

The experience would be different from simply watching a film. The user will be integrated with the fictional landscape (just as the fictional landscape will be integrated with the actual landscape in which the user lives) and will experience it as though he were in it. Moral education (indeed, any kind of education), though, will be problematic, at best, if the user is unable to correctly identify and interact with the virtual fictional beings that appear in his visual field. Suppose that Bradley Pearson appears in one's total experiential field. To the user, Pearson appears to be a real person: he is visible, his speech can be heard, he can be touched by the user. The authentic identification problem is two-fold. The first identification problem is not unique to fictional virtual beings – it involves any kind of virtual being, whether it is fictional or non-fictional. The problem is that the user must be able to distinguish real objects in his local environment from the virtual objects that appear in his total experiential field. If he fails to make the proper distinctions, his life might go quite badly.

The second identification problem is how virtual fictional beings are distinguished from virtual beings that are not fictional. The Bradley Pearson that appears in one's total experiential field is a virtual fictional being. Suppose another virtual being appears in one's total experiential field, but this being is not fictional. It might, for instance, be a virtual object that looks like a human being but functions to transmit information to the user. In that way, it is not a fictional being, since, it is not in a fiction of any kind. Suppose it looks like Bradley Pearson (as he is described by Iris Murdoch). If the user cannot distinguish the fictional Bradley Pearson from the virtual object that looks like Bradley Pearson, he will encounter serious problems in performing, among other things, everyday activities. The fictional Bradley Pearson might say it is raining outside, while the non-fictional Bradley Pearson might say it is not raining outside (if his function is to inform the user of weather conditions). For any sentence about one's local environment, a fictional virtual being might say it has one truth-value, while the non-fictional virtual being says it has the opposite truth-value. Sometimes the two will agree in the truth-values they assign to sentences. Where a user of augmented reality technology cannot distinguish fictional virtual beings from non-fictional virtual beings, he might be wrong about the truth-values of many sentences about the condition of his local environment.

Fictional entities modeled on real persons

A complication in the authentic identification problem arises when fictional virtual beings are modeled on real persons. Suppose we read a fictional story in which the central character is named "Bill Clinton." Suppose that in this story the character does some, but not all, of the things that the actual Bill Clinton did, as well as having some, but not all, of the physical and psychological characteristics of the actual Bill Clinton. Should we say that the fictional Clinton is a different entity from the actual Clinton or should we say that there is only one entity? Notice that we can make different sorts of assertions about the fictional and actual Clinton, depending

upon which of the two views we subscribe to. For instance, if we take it that there are two Bill Clintons, then we can say such things as that we dislike the actual Clinton, but like the fictional Clinton immensely. However, if there is only one Bill Clinton, then we cannot say that. Instead, we can say things such as that we really like Bill Clinton as he was described in the novel, but not as he was in actual life.

Which view is more natural: that there are two entities or that there is just a single entity? What depends upon the answer we provide will be both the kinds of dialogues we can carry out in augmented reality between actual and fictional beings and the kinds of assertions actual beings can make about fictional beings. If we take the one entity view, then when we observe the virtual Bill Clinton, we can truthfully say that he looks better as a virtual being than he now looks as an actual being.

Moreover, what we think we are seeing when we observe Bill Clinton in our augmented reality landscape will depend upon which view we adopt. If we take the one entity view, then when we see the virtual fictional Bill Clinton talking with his secretary of state, Madeleine Albright, we are actually looking at Bill Clinton. Similarly, when we look at videotapes of Clinton's impeachment proceedings, we are actually looking at Bill Clinton. However, if we take the two entities view, then when we look at the virtual fictional Bill Clinton, we are not looking at Bill Clinton, but at some fictional character who is named "Bill Clinton." When we look at the impeachment proceedings, we are looking at the actual Bill Clinton, unless the impeachment proceedings are part of the virtual fiction that appears in our total experiential field – in which case, we are not looking at Bill Clinton, but at some fictional character who is named "Bill Clinton." But, on the two entities view there is no connection between the two entities that would allow us to draw inferences about one based on what we observe about the other.

The question thus arises: when is there a need for a one entity view and when is there a need for a two entities view? Notice that if a two entities view is adopted, an extra referent for the name "Bill Clinton" has been created. Is the need determined by the intentions of the creator of the fiction? Or is it determined by the intentions (and perhaps other things) of the users of the virtual fictional being? For instance, suppose that it has been established that there is no actual person on whom James Bond is modeled. What happens, then, if Smith, in whose total experiential field the virtual fictional James Bond appears, decides that there is just one entity: the actual man who is named "James Bond" and who lives down the street from him? Is this a legitimate decision or can we fault it on the grounds that the James Bond who lives down the street from Smith is not the fictional James Bond who figures in his reality augmentation? Notice that if Smith adopts the one entity view, then there is nothing to stop anyone else from doing the same in similar situations involving the same or different fictional characters. In which case, there would no longer be any difference between being right or being wrong about fictional entities and their relation to actual entities, since any differences between actual and fictional entities would evaporate. I take this to be a powerful consideration against the view that we can simply decide whether we want to adopt a one entity or two entities view.

Fictional entities modeled on real persons: employing the two-entity view

Suppose that a virtual fictional entity is based on a real person. One can easily imagine the usefulness of virtual characters who resemble actual people who have close, if not intimate, relationships with the user. Smith goes to a strange city for a job that will last several months, while his spouse must remain in another part of the world during that period of time. Smith is subject to depression whenever he has feelings of loneliness, and is a devout man of faith who is, however, incapable of refusing the attentions and attractions of another woman. The reality

augmentation employed would have an image of his wife in his visual field at all times while he was awake.

It was earlier claimed that the one entity view is more natural than the two entity view. On this view, it is natural for Smith to say that he likes the way his wife is portrayed in the reality augmentation as well as in real life. On this view it is not natural for him to say that he does not like the reality augmentation wife, but he does like his real wife. However, there are occasions on which a two entities view is more natural than a one entity view. Certainly it is easy to imagine situations in which it is more natural to create a fictional character than to say that there is just one being, the actual person who is his wife. Obvious cases are where he encounters his actual wife, in which case he does not want to conflate the virtual wife with the actual wife. But there are other cases where a two entity view is more natural than a one entity view, even where Smith does not encounter his actual wife. Here is one such case: Smith witnesses his virtual wife sadistically pummel a tough-looking virtual character. An augmented reality containing this scenario is designed to discourage Smith from visiting a part of the new city in which he might encounter such people. It would be better that Smith distinguishes the virtual wife from his actual wife, and not think of his virtual wife that it is one way in which his actual wife is portrayed. The reason is fairly clear: it might easily happen that Smith makes fallacious inferences about his actual wife or that he comes to think of his actual wife in the same way that he thinks of his virtual wife. Notice that the one entity view might encourage making fallacious (context conflating) inferences.

But this raises a problem: are there any general principles by which it can be determined whether the one entity or the two entity view is most natural? It would appear there are no such principles, since it appears to be a contingent matter as to how someone will react to the virtual characters populating their reality augmentation. It is, admittedly, a complex contingent matter, since it partly depends upon the psychology of the user and that psychology can either remain the same under the reality augmentation or itself change under it in ways that are wholly unpredictable.

Levels of fiction and their relation to reality

It is commonplace that there are fictions within a fiction. In principle, there could be arbitrarily many levels of fiction within a fiction. This creates new difficulties for the cross-barrier communication problem, the foolish inference problem and the authentic identification problem. With the introduction of new levels of fiction, there are more ways in which contexts can be conflated and thus more ways in which the three problems for augmented reality involving fictional entities are exacerbated. There are numerous instances in fiction of a fiction within the fiction. One well-known example is Shakespeare's play *Hamlet*. In it, a fictional play called "The Murder of Gonzago" is staged – thus there is a play within the play. This raises various issues: Is there a fictional character Gonzago? Can we say, in speaking about the play *Hamlet*, that there are in it the fictional characters Hamlet and Gonzago? The answer is that we cannot. From outside the play, speaking about what happens in it, it is true that there is no such fictional character as Gonzago in the play. It is true that there is a fictional character Hamlet in the play.

The reason why it is false to say there is a fictional character named "Gonzago" in *Hamlet*, but true to say that there is a fictional character named "Hamlet" in it, is that Gonzago is a fictional fictional character and not a fictional character. Suppose we wish to speak within the play – that is, we wish to make assertions that have the truth-conditions of what is in the play. In which case, it is true that there is a real person named "Hamlet," but false that there is a real person named "Gonzago." Rather, it is true that there is a fictional character named "Gonzago." Now suppose we wish to make assertions that have the truth-conditions of "The Murder of Gonzago." That is,

we speak within the fiction within the fiction. If we do this, then we would say that there is no such fictional character as Gonzago. Rather, there is a real character Gonzago.

What about Hamlet from within the play within the play? Can we say that there is a fictional character named “Hamlet” in the play *Hamlet*? No, we cannot. The reason is obvious: From within the play “The Murder of Gonzago” there is no such fictional or real entity as Hamlet or even a play with the title “Hamlet.” From within “The Murder of Gonzago” there is no such fictional character or real person named “Hamlet.” Similarly, from within “The Murder of Gonzago,” there is no such real person named “President Obama.” What we can legitimately say “within the fiction within the fiction,” is a contested matter, since Shakespeare (in *Hamlet*) says almost nothing about what happens in the play “The Murder of Gonzago.”

To make matters perspicuous, we contrast assertions about Gonzago and about Hamlet that are each made within three different contexts:

- (i) In real life: It is true that there is no such fictional character as Gonzago.
- (ii) Inside the play: It is false that there is no such fictional character as Gonzago.
- (iii) Inside the play within the play: It is true that there is no such fictional character as Gonzago.

The difference between (i) and (iii) is that even though the assertion is true in both cases, it is true in (i) because Gonzago is a fictional fictional character, while it is true in (iii) because Gonzago is a real person. Now consider Hamlet:

- (iv) In real life: It is true that there is a fictional character Hamlet.
- (v) Inside the play: (i) It is false that there is a fictional character Hamlet. (ii) It is true that there is a real person Hamlet.
- (vi) Inside the play within the play: (i) It is false that there is a fictional character Hamlet. (ii) It is false that there is a real person Hamlet.

The difference between (v) and (vi) is that even though the first assertion – (i) in each – is false, it is false in (v) because Hamlet is a real person and it is false in (vi) because Hamlet is neither a real person, a fictional character, nor an iteration-of-fictions character.⁴

It makes no sense, however, to speak of utterances made within the play within the play, since in *Hamlet* the play within the play is only alluded to, but is not actually made into a work of fiction. If Shakespeare had made it into a work of fiction within *Hamlet*, then we might be able to provide identity conditions for the characters and events that would have occurred – had Shakespeare actually penned them – in “The Murder of Gonzago.” But without being made into a work of fiction, we cannot do that and so we cannot say, in either the context of reality or the context of within the play anything about the character Gonzago other than that Gonzago is a fictional character (in the context of *Hamlet*) and a fictional fictional character (in the context of reality). There are interesting cases in fiction of fictions within the fiction that are substantial. For instance, in Fritz Lang’s film *The Woman in the Window*, the fiction within the fiction occupies most of the film. Importantly, the character played by Edward G. Robinson is the same in both the fiction and the fiction within the fiction. However, from within the fiction within the fiction, one cannot say that it is true that that character is a fiction. Rather, from inside the fiction within

the fiction, that character does not exist, in any iteration of levels of fiction. Rather, one can only say within the fiction within the fiction that he is an actual person.

The basic problem that the existence of different levels of fiction creates for augmented reality is that unless the utterances made within the reality part of augmented reality are subscribed, there is no way of correctly determining whether they are true or false. This creates a difficulty for any user of augmented reality technology whose augmented mental life is replete with virtual fictional entities, each of which occupies some level of fiction, though not the same level for each. Simply hearing words uttered by these beings will not help in determining which level of reality they occupy, and thus not help in determining the truth-conditions of those utterances.

Images as representing information and as representing fictional entities

Suppose a virtual fictional character, such as Mickey Mouse, appears on your computer screen. Mickey announces to you the time of day whenever you log on. It would be a mistake to say that the fictional character Mickey Mouse is communicating with you. Rather, there is a virtual pictorial representation of a fictional character on your screen and there are sounds that appear to come from the mouth of that representation. These sounds convey information about the actual world. They are not words in the fictional world of Mickey Mouse, since the pictorial representation of Mickey Mouse is not intended to be a piece of fiction. Rather, the spoken words and the picture are symbols that are intended to convey information about the actual world. The image on the computer screen is the pictorial representation of the fictional character Mickey Mouse. Of course, it could also be the case that the virtual pictorial representation of Mickey Mouse is intended to be the representation of the fictional Mickey Mouse and that the utterances made are intended to be the fictional Mickey telling, within the fictional world, a being outside that world what that being wants to know about the correct time of day in the actual world that being inhabits. It can be part of a fiction that the characters in it acknowledge they are fictions and that they acknowledge that they are talking to the reader of the fiction. If that is so, then from the point of view within the fiction, beings outside the fiction who are reading the fiction are fictional fictional characters (because Mickey takes himself to be a fictional character in the fiction).⁵ The existence of an “outside world” is just a fiction within the fiction. The owner of the computer who hears the correct time uttered by Mickey Mouse, is a fiction from the point of view of the fiction. But this will be so only if it is the intention of the designer of the virtual pictorial representation of Mickey Mouse for the computer that she pretend that Mickey Mouse is pretending to provide the time for a fictional character who needs, in his fictional world, to know the time that is, within that fictional world, the correct time.

What shall we take it to be, when we confront a virtual pictorial representation of Mickey Mouse on our computer screen? Just as important, does it matter, in any way, what we take it to be? Is it a virtual fictional Mickey engaged in a fiction or is it a virtual non-fictional Mickey that conveys information about the actual world? If the virtual fictional Mickey appears in Smith’s augmented reality, is it essential that Smith be able to tell that Mickey is fictional and not non-fictional? If Smith cannot make the distinction, then he will succumb to any of the three problems for augmented reality involving fictional entities. Where the correct answer to the question of whether a virtual being is either fictional or non-fictional depends on the intentions of the designer of the augmented reality technology, there may be no determinate answer, since the intentions of the designer might themselves be indeterminate. (That one’s intentions – with respect to what one means by the words one uses – might be indeterminate is argued in Kripke 1982.)

Virtual pictorial beings without designers in augmented reality

Suppose that the virtual pictorial non-fictional beings inhabiting someone's augmented reality are randomly chosen from a catalog by a computer program whose designer intends only that the subroutine that chooses them from a catalog does so randomly. The subroutine that creates the virtual pictorial non-fictional beings has no designer intentions. There is sophisticated software that enables the virtual pictorial non-fictional beings to mesh with the visual features in one's visual field in such a way that the being does not replace a visual feature that is important for navigating one's way in the actual world without injury or without losing or compromising important information necessary for one's projects.

There is also no intention on the part of the designer that the virtual pictorial non-fictional beings be anything more than that – there is no intention that they are fictional characters and that their activities are a pretense that is part of a fictional world. If so, human verbal interactions with virtual pictorial non-fictional beings are just like human interactions with a clock or with an electronic audio recorder. We see the time the clock registers and do not take the registered time to be part of a fictional world.⁶ We hear sounds coming from the electronic audio recorder and, unless we know that the sounds are from a fiction, such as actors speaking as they play fictional characters, we do not take them to be utterances of fictional characters, but do take them to be utterances of actual people.

Interactions with virtual pictorial non-fictional beings that inhabit the landscape of our augmented reality are like interactions with real things in the real world. The software that manipulates the position of such beings in order to either prevent an accident (such as when they mask a large rock on the pavement that we might trip over as we walk) or to accommodate our plans and projects does not perform the manipulations as part of a pretense. Rather, the software receives information about the actual world and makes adjustments in the activities of the beings on the basis of that information. Once again, how would we tell when a virtual pictorial being is non-fictional and without a designer rather than a virtual pictorial fictional being or a virtual pictorial non-fictional being with a designer? If a user of augmented reality cannot make these distinctions, then he will succumb to the three problems of augmented reality involving fictional entities and to an additional problem. If the user thinks that the virtual pictorial being is in his total experiential field owing to designer intentions, then he might falsely infer that there is a purpose for the being, falsely infer a purpose, and then act as though that purpose is fact.

Computational simulation of human intelligence and fictional entities

Suppose that the software technology capable of achieving augmented reality with fictional entities can simulate human intelligence. However, if it is the case that a human user of this technology will make fallacious inferences if he is not able to distinguish between fictional and real beings with respect to their utterances, or is not able to distinguish between fictional and real beings with respect to their properties (such as height, weight, etc.), then that technology will not be capable of providing a safe environment for the user, even if it is sophisticated enough to computationally simulate human intelligence. If we do not know the intentions of the designer of the technology with respect to the ontological status of the fictional beings that occur within augmented reality, then we will not know whether we are correctly reasoning about them. There are various possibilities for error: we take virtual pictorial beings to be fictions, while it is the intentions of the designer that they are non-fictional. If so, we reason incorrectly about them. Or we take them to be non-fictional, while it is the intentions of the designer for them to be fictions. If so, we reason incorrectly about them.

The intentions of a designer of the technology necessary for achieving augmented reality might be wholly indeterminate. Or it might be a matter of fact what they are, but we do not have access to those facts. In either case, if we do not know the intentions of the designer, then we will not know whether the virtual pictorial beings are non-fictional and whose utterances convey information about the actual world or are fictional entities and whose utterances are pretenses that are part of a fictional world it was the intention of the designer to create. Ordinarily, we would not worry about this kind of skeptical problem. However, where the virtual pictorial beings are part of augmented reality, it will matter, since we are using the information conveyed in their utterances to successfully navigate in the actual world.⁷

Facts, fictions and forecasts

Facts within a fiction can be fictions outside the fiction, though this is not necessarily so. Consider a science-fiction story written in 2011 in which a scientist in the year 2421 creates a time travel machine called “The Slingshot.” Within the fiction, it is a fact that “The Slingshot” was created in 2421. Outside the fiction, it is a fiction that “The Slingshot” was created in 2421. Outside the fiction, we can say that inside the fiction the real entity “The Slingshot” was created in 2421. We can also say, outside the fiction, that the fictional entity “The Slingshot” was created by its author in 2011. Whether something is a fact or a fiction depends upon the context in which the assertions are made. Indeed, predicates are also context-sensitive. Kripke’s example of a contextualized predicate is “created” (Kripke forthcoming 2012, 7-8 Lecture 4). In saying that Frankenstein’s monster was created whenever the novel says that Frankenstein created it, the predicate “created” is relativized to the context “fictionally.” However, in saying that Frankenstein’s monster was created in 1818, the predicate “created” is relativized to outside the novel (Kripke 2011; Kripke uses the phrase “out-and-out” for the context in which attributions about what happens inside a fiction are made outside the fiction). What is a fact within the fiction is a fiction outside the fiction. However, whether something is a fact or a fiction will depend upon a context. That this is so creates problems for making factual assertions in augmented reality.

Suppose that information needs to be communicated within different parts of an augmented reality in which some parts are real and some parts are fictional. We can easily distinguish between the fiction we see in a film or read in a book from the surrounding reality. When, however, the real and the fictional form an integrated total experiential field, there may be no means of correctly identifying the contexts of parts of it. That is, it might be infeasible to correctly identify fictional and real beings and thus, to identify utterances within a fiction and utterances within reality. If a fictional entity utters a sentence that is heard by a real entity, then whether it is communicated to the real entity and whether it is true or false will depend upon the context. However, if no context can be specified, it will be impossible to determine whether it is true or false and whether it is communicated to a real entity. (For the most part, specification of a context in asserting facts is not necessary where there is a single reality.) Suppose that the information which is communicated is the birth date of Frankenstein’s monster. Whether it is a fact or a falsehood (that is, a fiction) will depend upon the context in which it is asserted. But if the context is left unspecified, there is no transmission of information, since the assertion might be true or it might be false. Since we have no more reason to think it is true than we have to think it is false, we are maximally ignorant as to its truth-value and thus have not acquired any information. It is not an assertion on which we can act, make a decision, or solve a problem.⁸

However, care must be employed in describing the context, since a true assertion made within a fiction is taken to be actual within it. Thus within a fiction the context is that of actuality: the actuality from within the fiction. If the utterances of fictional beings in augmented reality are indexed by “actuality,” those indices must be changed when the utterance is heard by a real being

– such as the user of augmented reality technology. The problem, though, is that it is only the utterances of fictional beings that require this procedure. It is not required when the virtual pictorial being is non-fictional and is used to convey information to the user. Thus it must be transparent when a communication comes from a fictional entity and when it comes from either a virtual pictorial non-fictional entity or an actual entity. If Smith hears a message from his wife that she is ready to leave for the movies, that information does not require redescribing the context. But if a virtual pictorial fictional entity says that she is ready to leave for the movies, that information does require redescribing the context. However, there may be circumstances under which it is difficult to determine whether information is transmitted from an actual or from a fictional entity, since the simulation of the actual entity within the fiction is so authentic that the user is fooled into thinking it is real.⁹

If we use information acquired in augmented reality to make forecasts, we had better be ready to make mistakes, since there may be circumstances under which we will not be able to distinguish between virtual fictional beings, virtual non-fictional beings and real beings. We could have a false belief in various ways. One way is to not be able to distinguish the three kinds of beings. Another way is to think (falsely) that a fictional being can communicate information to oneself – the user of augmented reality. A third way is to have a false philosophical theory of fictional entities. In all three cases, the upshot is that the user will have false beliefs and these false beliefs will result in using information in ways that will jeopardize his well-being. Worse yet, where virtual fictional, virtual non-fictional and real entities are observationally indistinguishable within the framework of the total experiential field of the user of augmented reality, the user of augmented reality technology cannot tell the entities apart by merely looking. In that case, natural predicates and fictional predicates might be entrenched in the same way in the belief formation mechanism of the user. If so, the user will form bizarre hybrid inferences, such as statistical generalizations that combine natural and fictional predicates. Forecasts based on such hybrid statistical generalizations will combine fictional and factual without distinguishing between them. Beliefs about the world will be fixed where the evidence might be purely fictional and beliefs about fictions (where a user has such beliefs) might be fixed by facts (of the real world). Where forecasts are based on a fusion of facts and fictions (that is, facts within the fiction), disaster may follow the user.

Hybrid inferences might, though, suppress the entrenchment of grue and grue-like predicates. Goodman (1983) defined disjunctive grue and grue-like predicates, but we can also define conjunctive grue and grue-like predicates. Grue, defined conjunctively, is: an object is grue if and only if it is green when observed by at least one person and blue when it is not observed by any person.¹⁰ Imagine fictions in which there are always fictional observers making observations. If so, then perhaps we do not have to worry that it will ever be the case that any objects that are green when observed will turn blue when unobserved. Since a part of ourselves – though fictional, and assuming we are persons – is making an observation, we are making the observation. Defeating grue by fiction? That is stranger than fiction. Such oddities are part and parcel of what might be in store for us when augmented reality involving virtual fictional and non-fictional beings becomes technologically feasible.

Notes

1. Given various definitions of health, what is an enhancement on one definition is a treatment on another. This opens up the possibility of creating pairs of cases that reverse when definitions of health are changed. On one definition of health (say that proposed by Normal Daniels), what is (unfairly) an enhancement for A is treatment for B and what is treatment for C is also treatment for D. Adopting a different definition of health, what is treatment for B is also treatment for A,

but what is treatment for D is (unfairly) an enhancement for C. I explore the ways in which definitions of health determine how we draw the treatment-enhancement distinction in a forthcoming paper.

2. If someone – a real person – dresses up as Hamlet, that is obviously different from the fictional character Hamlet appearing in one’s visual field. It is a real person who adopts a disguise or costume that is modeled on the fiction.

3. Of course, you could be thinking both about the image of Hamlet *and* also about the fictional character Hamlet, but in that case, you would be thinking about two different things, and not just about one thing.

4. Inside the play “The Murder of Gonzago,” there is no such entity named “Hamlet.” Thus there are no statements inside the play about which one can say they are true or false. This raises the issue of existential presuppositions in the ascription of truth-values to statements about objects that do not exist, but the issue – that a statement uttered within “The Murder of Gonzago” might have no truth-value or a truth-value other than true or false – can be side-stepped here without affecting the main points about levels of fiction.

5. Where part of the pretense within a fiction is that the fiction is itself a fiction, new complexities arise. The complexity of embeddings of pretense fictions within a fiction could make it extraordinarily difficult to assign levels of fiction indices to parts of one’s total experiential field.

6. That is not to say that we could not imagine – or pretend – that it is part of a fictional world. Another kind of interaction we might have with entities in our total experiential field is that we pretend that they are fictions, even though they are not. That is, we can take our interactions with them to be a pretense. Needless to say, this kind of interaction introduces enormous difficulties for the task of assigning levels of fiction indices to parts of one’s total experiential field.

7. Suppose that – given software sophisticated enough to do this – the fictional entities occupying part of your total experiential field create new fictional entities. This is, of course, reminiscent of “The Murder of Gonzago.” Whose intentions do we use to determine whether the new entities we experience are indeed fictional?

8. This difficulty can also be considered a form of skepticism, since in cases where we cannot know the truth-value of our knowledge-claims, we cannot have knowledge.

9. This is a form of skepticism. It might be blocked by arguing that fictional entities are no less real than real entities, since all are computations and so all are equally real. Recently, David Chalmers has argued that the true metaphysics of the world might be computations and thus such a fact about the world would undercut brain-in-a-vat skepticism (in his paper “The Matrix as Metaphysics,” available on his webpage). The idea is that all of the images of the world that are fed into the brain-in-the-vat by evil vat scientists are computations and thus real. Similarly, it might be argued that all of the virtual fictional entities in augmented reality are just as real as the virtual non-fictional entities and the real entities. There are several ways in which it can be argued that the metaphysics of the world is computations. In a forthcoming paper I argue that Chalmers’ argument for the computational nature of reality (i) would imply that there is a collapse of the hierarchical type structure of intensional objects into a single level and (ii) is in direct conflict with Kripke’s views on fictional entities and the logic of non-existence.

10. The conjunctive definition of grue is my own invention. Goodman does not define grue conjunctively and has, though not in print, disavowed any interpretation of grue that takes the grue problem to be a form of skepticism. Obviously, the conjunctive interpretation of grue offered in the text is a form of skepticism in addition to its role in revealing an infirmity in confirmation theory.

References

Gurwitsch, Aron. 1964. *The field of consciousness*. Pittsburgh: Duquesne University Press.

Goodman, Nelson. 1983. *Fact, fiction, and forecast*. 4th ed. Cambridge: Harvard University Press.

Hintikka, Jaakko. 1963. Cogito, ergo sum as an inference and a performance. *Philosophical Review* Volume LXXII, Number 4: 487-96.

Kripke, Saul. 1974. Reply to David Kaplan. Unpublished manuscript.

Kripke, Saul. 1980. *Naming and necessity*. Cambridge, MA: Harvard University Press.

Kripke, Saul. 1982. *Wittgenstein on rules and private language*. Cambridge, MA: Harvard University Press.

Kripke, Saul. 2011. Vacuous names and fictional entities, in *Philosophical troubles: Collected works*. Volume 1. 52-74. New York: Oxford University Press.

Kripke, Saul. Forthcoming 2012. *Reference and existence: The 1973 John Locke Lectures*. Collected Works, Volume 2. Oxford University Press.

DR. RUPNATHU (DR. RUPAK NATH)