

APA Newsletters

NEWSLETTER ON PHILOSOPHY AND COMPUTERS

Volume 09, Number 1

Fall 2009

FROM THE EDITOR, PETER BOLTUC

FROM THE CHAIR, MICHAEL BYRON

NEW AND NOTEWORTHY: A CENTRAL APA INVITATION

ARTICLES

Featured Article

RAYMOND TURNER

“The Meaning of Programming Languages”

GREGORY CHAITIN

“Leibniz, Complexity, and Incompleteness”

AARON SLOMAN

“Architecture-Based Motivation vs. Reward-Based Motivation”


DISCUSSION 1: ON ROBOT CONSCIOUSNESS

STAN FRANKLIN, BERNARD J. BAARS, AND UMA RAMAMURTHY

“Robots Need Conscious Perception:
A Reply to Aleksander and Haikonen”

PENTTI O. A. HAIKONEN

“Conscious Perception Missing. A Reply to
Franklin, Baars, and Ramamurthy”



ONTOLOGICAL STATUS OF WEB-BASED OBJECTS

DAVID LEECH ANDERSON

“A Semantics for Virtual Environments and the Ontological Status of Virtual Objects”

ROBERT ARP

“Realism and Antirealism in Informatics Ontologies”

DISCUSSION 2: ON FLORIDI

KEN HEROLD

“A Response to Barker”

JOHN BARKER

“Reply to Herold”

DISCUSSION 3: ON LOPES

GRANT TAVINOR

“Videogames, Interactivity, and Art”

ONLINE EDUCATION

MARGARET A. CROUCH

“Gender and Online Education”

H. E. BABER

“Women Don’t Blog”

BOOK REVIEW

Christian Fuchs: *Social Networking Sites and the Surveillance Society. A Critical Case Study of the Usage of studiVZ, Facebook, and MySpace by Students in Salzburg in the Context of Electronic Surveillance*

REVIEWED BY SANDOVAL MARISOL AND THOMAS ALLMER

SYLLABUS DISCUSSION

AARON SLOMAN

“Teaching AI and Philosophy at School?”

CALL FOR PAPERS

Call for Papers with *Ethics Information Technology* on
“The Case of e-Trust: A New Ethical Challenge”

Baars, Bernard J. 2002. The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Science* 6: 47-52.

Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., et al. 2006. High gamma power is phase-locked to theta oscillations in human neocortex. *Science* 313: 1626-28.

Franklin, S. 2003. IDA: a conscious artifact? *Journal of Consciousness Studies* 10: 47-66.

Franklin, S. 2005. A “consciousness” based architecture for a functioning mind. In *Visions of Mind*, edited by Darryl N. Davis. 149-75. Hershey, PA: Information Science Publishing.

Gaillard, R., Dehaene, S., Adam, C., et al. 2009. Converging intracranial markers of conscious access. *PLoS Biology* 7(3): e1000061.

Haikonen, Pentti O.A. 2009. Slippery steps towards phenomenally conscious robots. *APA Newsletter on Philosophy and Computers* 08:2.

Jensen, O., & Colgin, L. L. 2007. Cross-frequency coupling between neuronal oscillations. *Trends in Cognitive Sciences* 11(7): 267-69.

Massimini, M., Ferrarelli, F., Huber, R., et al. 2005. Breakdown of cortical effective connectivity during sleep. *Science* 309: 2228-32.

Mitchell, M. 1993. *Analogy-making as Perception*. Cambridge MA: The MIT Press.

Uchida, N., Kepecs, A., and Mainen, Zachary F. 2006. Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nature Reviews Neuroscience* 7: 485-91.

van Berkum, J. J. A. 2006. *Discourse and the Brain*. Paper presented at the 5th Forum of European Neurosciences: The Federation of European Neuroscience Societies (FENS), Vienna, Austria.

Willis, J., & Todorov, A. 2006. First impressions: making up your mind after a 100-Ms exposure to a face. *Psychological Science* 17: 592-99.

control structures) implement the essential elements of Global Workspace Theory.” This is not a modest claim at all.

Personally I would be quite happy if it could be shown that my cognitive model (Haikonen 2003, 2007) had captured some elements of human cognition (I trust it has), but I would not dare to claim that the brain implements my model, even in the unlikely case where my model would turn out to be a perfect model of the brain. The human brain and mind are a little bit older constructions than the Global Workspace Theory and have evolved without any knowledge of the same. It seems that here FBR have switched the role of a natural object and its man-made model. No natural object or system is based on man-made models or blueprints. To claim the opposite is to nominate oneself as the Creator. However, we are free to find Nature’s principles and implement those in our own designs.

References

Franklin, Stan, Baars, Bernard J., Ramamurthy, Uma. 2009. A phenomenally conscious robot. *APA Newsletter on Philosophy and Computers* 08:2.

Boltuc, Peter. 2009. The philosophical issue in machine consciousness. *International Journal of Machine Consciousness* 1: 155-76.

Haikonen, Pentti O.A. 2009. Slippery steps towards phenomenally conscious robots. *APA Newsletter on Philosophy and Computers* 08:2.

Haikonen, Pentti O.A. 2003. *The Cognitive Approach to Conscious Machines*. UK: Imprint Academic.

Haikonen, Pentti O.A. 2007. *Robot Brains: Circuits and Systems for Conscious Machines*. UK: Wiley & Sons.

Conscious Perception Missing. A Reply to Franklin, Baars, and Ramamurthy

Pentti O. A. Haikonen

University of Illinois at Springfield

Franklin, Baars, and Ramamurthy kindly clarify their position in response to my critique (Haikonen 2009). I fully recognize the important and pioneering work that Franklin, Baars, and Ramamurthy have done in the field of artificial cognition and my critique should not be construed to diminish the value of that work in any way.

However, in the good tradition of philosophical debate I would like to point out the following. There seems to be nothing in the writings of FBR (or anybody else’s) that would explain how the running of any computer program could evoke qualia and subjective feelings in the executing machine. On the other hand, it is obvious that computer programs can simulate various feelings including pain and pleasure via their functional consequences. The presence of such consequences does not, however, prove that the computer would actually feel something or be conscious in the hard sense (h-consciousness, see Boltuc 2009). It may well be that the phenomenal aspects of consciousness are beyond the capacity of computer programs and may be present only in some hard-wired perceptive and reactive systems.

FBR have done excellent work in the development of the LIDA agent that they call functionally conscious. Based on that they wish to define functional consciousness as the process that implements the Global Workspace Theory and the LIDA Cognitive Cycle (Franklin, Baars, and Ramamurthy 2009) a notion that I criticized in my previous response. The concept “functional consciousness” is doubtful, but even so, it should not be hijacked to apply to one specific cognitive model only. In their current response FBR wish to go even further. They state: “There is a sizable and growing body of evidence from cognitive science and neuroscience that human minds (their

ONTOLOGICAL STATUS OF WEB-BASED OBJECTS

A Semantics for Virtual Environments and the Ontological Status of Virtual Objects

David Leech Anderson

Illinois State University

Abstract

Virtual environments engage millions of people and billions of dollars each year. What is the ontological status of the virtual objects that populate those environments? An adequate answer to that question requires a developed semantics for virtual environments. The truth-conditions must be identified for “tree”-sentences when uttered by speakers immersed in a virtual environment (VE). It will be argued that statements about virtual objects have truth-conditions roughly comparable to the verificationist conditions popular amongst some contemporary antirealists. This does not mean that the virtual objects lack ontological standing. There is an important sense in which virtual objects are no less real for being mind-dependent.

Introduction

What is the ontological status of the virtual objects that populate the burgeoning virtual worlds that reside on the Internet? *Second Life* is a virtual world comprised not only of objects like tables, chairs, trees, and fountains, but large landmasses that have cities with expansive real estate developments. The people who frequent this virtual environment (by animating virtual bodies known as “avatars”) not only prize these virtual objects in some Platonic way, they place a commercial value on them by paying cold, hard cash. Anshe Chung (a.k.a. Ailin

Graef) became a *cause célèbre* and made it on the cover of *Business Week* magazine when the value of her combined virtual holdings in *Second Life* exceeded \$1,000,000.¹ Not one million in virtual “Linden” dollars (the currency within *Second Life*), but one million dollars U.S.

World of Warcraft, with over 11 million paying monthly subscribers, is another massively multiplayer online role-playing game (MMORPG) that is played within a virtual environment. The virtual objects that populate games like this are bought and sold earning their creators hundreds of millions of dollars each year.² One might begin a discussion about the ontological status of such virtual objects by invoking a famous claim advanced by Ian Hacking, a party to the realism-instrumentalism debate on the ontological status of theoretical entities (like electrons). Hacking reports on a conversation he had with a working physicist. He recounts:

Now how does one alter the charge on the niobium ball? “Well, at that stage,” said my friend, “we spray it with positrons to increase the charge or with electrons to decrease the charge.” From that day forth I’ve been a scientific realist. So far as I’m concerned, if you can spray them then they are real.³

In the same spirit, one might be tempted to say: “If you can buy them and sell them for hundreds of millions of dollars then they are real.” Even so, merely acknowledging the reality of virtual objects gets us no distance to understanding in what that reality consists. The first step must be to determine what counts as a “virtual object” and then we can ask where (if at all) in our ontological hierarchy they are properly placed.

Virtual environments and virtual objects

From the outset, the meaning of “virtual object” (VO) will be restricted so as to exclude some digital entities that are so described in discussions about Internet commerce. The day of Michael Jackson’s memorial service, Facebook gave away 800,000 copies of what was described as a “special commemorative virtual gift”—a graphical representation of Michael’s white sequined glove. In this discussion, graphical images like the white sequined glove will not be treated as a virtual object, because a Facebook page does not qualify as a genuine “virtual environment.” It lacks nothing in its *virtuality*, but for the purposes of this discussion it isn’t rich enough to constitute an *environment*, which will be defined as follows:

E = an *environment* is a dynamic space-time region (with a minimum of two but typically three spatial dimensions) populated by objects that bear spatial and temporal relations to one another.⁴

The environments which we typically inhabit are physical in nature.

PE = a *physical environment* is a dynamic space-time region that consists of objects that bear spatial and temporal relations to one another and objects whose identity conditions include intrinsic, non-relational properties that exist independent of the present cognitive states (thinking, believing, experiencing, etc.) of any cognitive agent.

An MMORPG also constitutes an environment. It provides objects to serve as the target for agents’ intentional states and a stage upon which human actions (buying, killing, lying, sharing, etc.) become intelligible and even morally evaluable. We will describe these environments as “virtual” and I propose that we define them as follows:

VE = a *virtual environment* is a dynamic space-time region that consists of objects that bear spatial and temporal relations to one another and whose identity conditions supervene on the actual (or possible) sensory and cognitive

states of the agents who inhabit that environment.

Virtual environments are populated with virtual objects.

VO = a *virtual object* is an empirically detectable, intersubjectively stable, publically accessible entity that can be identified and re-identified over a sufficiently long run and whose identity conditions supervene on the actual (or possible) sensory and cognitive states of the agents inhabiting the associated VE.

The reader will notice that the definitions of VE and VO make their very existence dependent upon the cognitive agents who engage with them. If a plague wiped out the entire human race, and there no longer existed agents for whom the virtual tree could be an object of experience, then the tree would cease to exist.

An opponent will object that this consequence is unreasonable. Virtual environments (VE’s) need not be these ephemeral things that blink out of existence when the cognitive agents who previously inhabited them cease to exist. VO’s could, instead, be defined so as to be constituted by the physical systems that are causally responsible for generating users’ experiences of them. In that case, even if all humans died in a plague, so long as the machinery kept running, the VO’s would continue to exist.

Those sympathetic to this latter interpretation of VO’s will object that my previous definition attempts to settle, by stipulation, what should be a central controversy of this paper. This is a fair objection. I concede that my definitions of VE and VO do require a convincing argument—an argument that will be offered below. Laying these admittedly contentious definitions on the table at the outset will, however, streamline the discussion and aid in the explication of the view. In the end the reader must judge whether the definitions are well motivated.

Virtual Environments (VE’s) in the history of philosophy

By making virtual environments dependent upon the cognitive activity of the agents inhabiting them, I am quite deliberately invoking the idealist / antirealist philosophical traditions and the various alternatives they offered to a traditional realist account of the nature of the external world. Consider how the language of VE’s and PE’s might be used to describe debates over realism from Descartes to the present. We might recast Descartes’ methodological skepticism by considering the possibility that while I am having the experiences of a tree in the quad (VO), there may not exist a physical tree (PO) that answers to it. We could then imagine Berkeley arguing that the very concept of a material tree (PO) is incoherent and all that normal people mean when speaking of a “tree” is the tree-as-experienced (the VO). Michael Dummett’s “language acquisition argument” will translate into the claim that there is no coherent account of how humans could learn to understand a language with realist truth-conditions (which requires asserting the existence of PO’s) and so the correct semantics for natural languages must be verificationist (which requires asserting the existence of VO’s only).⁵ Finally, Hilary Putnam’s “Brains in vat” argument can be seen as an attempt to show that even if one begins with externalist assumptions, one will utterly fail in one’s attempt to raise the Cartesian specter of radical skepticism by considering the case where I am a disembodied brain floating in a vat of nutrients stimulated by a computer (experiencing VO’s but not PO’s). Ironically, knowing that I am not a brain in a vat does *not* confirm the truth of realism, according to Putnam.⁶ On the contrary, if I can’t raise the specter of radical skepticism—because there is no genuine possibility that I could be in a VE that is not being caused by a corresponding

PE—then the very distinction between VE’s and PE’s collapses, as most antirealists insist, and we are left with the incoherence of metaphysical realism.⁷

Admittedly, I am playing fast-and-loose with these august traditions as I reformulate them in terms of VE’s and VO’s. But I do hope that in spite of any quibbles one might have about my reading of history, the general point can still be made. The territory we are exploring is not unique to the twenty-first-century world of MMORPG’s. And the decisions we make about the proper analysis of VE’s might conceivably speak to or even commit one to certain positions in much broader areas of philosophy.

One more comment before we progress. The reader should be cautioned not to conclude that embracing a verificationist semantics for virtual objects in any way counts against the truth of metaphysical realism regarding the external world. Quite the contrary is actually the case, or so I would argue if space permitted.

Ontology with semantics

Our ultimate goal is to determine the ontological status of virtual objects. One cannot determine the ontological status of a particular virtual tree, however, unless one first determines what the virtual tree *is*. But that is ultimately related to the semantic question, What are the truth-conditions of the sentence, “There is tree” when spoken about the virtual tree in the virtual environment? The point is *not* that you can ignore metaphysics and simply do semantics. Certainly not. The point is that reality is too metaphysically rich; there are simply too many realities that are *prima facie* candidates for being the referent of the virtual “tree” in question. Nothing will be accomplished if one gives the most elegant ontological account of a phenomenon that one takes to be a virtual tree, if everyone in your audience insists that the phenomenon that you analyzed is simply *not* a “virtual tree” given what everyone else means by those words. The war is only won if you win both the semantic battle and the ontological battle. There is no avoiding the semantic question. We want to know the ontological nature of virtual trees; but to answer that question, we must also learn what the truth-conditions are for “tree”-sentences uttered within virtual environments. It is to both questions that we now turn.

A slippery slope argument

When native English speakers are immersed in an MMORPG and utter a “tree”-statement, the language they are speaking (*whatever* language that turns out to be) we will call “V-English.” I have proposed that the truth-value of such “tree”-sentences is not sensitive to the state of any physical object and is instead determined solely by how things stand with respect to the cognitive states of the inhabitants of the game. Thus, “tree”-sentences will be true on this account if all or most inhabitants of the game are having cognitive experiences constitutive of an “empirically detectable, intersubjectively stable, publically accessible” tree.

An alternative to this view is one motivated by the currently very popular position of semantic externalism. On one version of this view, a word refers to whatever it is that “causally regulates” the use of that term. The referent of the term is whatever it is that lies at the *end* of the causal chain that ultimately causes the speaker to utter sentences like, “Look, there is a tree.” For a speaker immersed in a particular MMORPG, there are a number of candidates that might fit that description. Let’s consider a few. Each of the numbered conditions that follow is a candidate for being that physical state-of-affairs that *causally regulates* the use of the term, “tree,” when uttered by a person immersed in a VE. When identified, that condition will be the referent for the term, “tree,” in V-English. The first candidate is:

1. *States of the server hardware*: The throwing of the electrical (on-off) switches on the VE-server that implements the “tree”-subroutine in the server software.

This has been a popular choice for an externalist referent of “tree”-statements as uttered by “brains in a vat” according to Putnam’s thought experiment,⁸ and at first blush seems equally promising here. The software alone is hard to target because it is an abstract entity, not a physical object. The switches alone, outside of the context provided by the “tree”-software, don’t capture the continuity of the tree through time. This articulation attempts to capture the best of both worlds, embracing both the server’s hardware and software. But this option is susceptible to a counterexample.

Assume that a bug is detected in the VE-server software of a famous MMORPG right in the middle of a well-publicized national contest being waged live and online between just two contestants. The two will soon be shooting (virtual) arrows into a tree and the bug can be expected to produce a malfunction. In order to prevent this eventuality, the programmer has a plan. The programmer knows exactly what signals (plus TCP/IP – Internet protocols) *would* be sent out over the Internet to the two contestants’ PC’s *if* the server were functioning properly. Imagine, as well, that the programmer has the ability (it doesn’t matter how) to interrupt the stream of defective instructions whenever they are sent by the software bug, and to send instead packets of instructions over the Internet that produce the proper “tree”-effects on the users’ computers. (Consider the programmers actions here on analogy with the actions of God in correlating the actions of minds and bodies according to “Occasionalist” theories of mind-body interaction.)

In this case—where the programmer is ensuring that inhabitants of the VE will continue to experience an “empirically detectable, intersubjectively stable, publically accessible” tree—what do we say about the truth-value of the sentence, “The tree was struck by an arrow” when uttered within the MMORPG? It seems only reasonable to say that it is true. The quick-thinking, spontaneous actions of the programmer preserved the existence of the tree within the VE. However, if the V-English word, “tree,” refers to condition 1. identified above, then the sentence must be false. Because of the software bug, the “tree”-routine is no longer being run. But “false” is the wrong result. As a matter of fact, it turns out neither hardware nor software is either a necessary or a sufficient condition for the existence of the virtual tree. Happily, the case itself suggests a second (PO) candidate for the referent of the word “tree”:

2. *Signals propagated over the Internet*: The carrying of the instructions plus TCP/IP Internet protocols that propagate the “tree-generating” instructions sent to the personal computers of all agents inhabiting the VE.

The reader herself, however, can probably generate a counterexample to this proposal. Instructions sent by TCP/IP is again just a half-way house. It is only a means for delivering instructions to the PC’s of each participating agent. There are any number of methods that might accomplish this, including tens of thousands of employees scattered around the world using all manner of quirky occasionalist methods for getting signals to all relevant computers causing the proper pixels to light up and forming the 3D image of a tree. So long as the result is a stable, intersubjectively consistent, genuinely public, virtual tree—“tree”-statements uttered in V-English will still come out true, even when condition 2. is lacking.

But now we are on a slippery slope. “Pixels lighting up on computer monitors around the world” is no more the proper “end” of the causal chain than any of the previous ones. Pick

any condition on the list. It is possible for *that* condition not to hold but so long as the next condition down on the list *does* hold, then the existence of the virtual tree will be preserved.

3. *The hardware in every user's PC*: The aggregate throwing of all appropriate electrical (on-off) switches in the machine hardware that implements the "tree"-subroutine on *all* the personal computers of *all* the agents inhabiting the VE.

4. *The monitor illumination in every user's PC*: The synchronized illumination of pixels on all the computer monitors of all agents so as to create an intersubjectively consistent 3D public tree.

5. *Retinal stimulation in every user's eyeball*: The proper stimulation of the retinas in all agents' eyes so as to create an intersubjectively consistent 3D public tree.

6. *Visual cortex stimulation in every user's brain*: The excitation of the proper areas of the visual cortex (V1–V5) in all agents' brains so as to create an intersubjectively consistent 3D public tree.

Each of the conditions is neither necessary nor sufficient for the existence of the virtual tree. It doesn't matter *how* you accomplish the task of bringing about the conditions described by VE and VO. It doesn't matter what physical system (or mental or spiritual system for that matter) is used along the way, all that matters is that you produce the final effect that literally *constitutes* the existence of the virtual tree. And that final effect inevitably is⁹:

7. *The intersubjectively coordinated, conscious experience of all users*: The proper production of an empirically detectable, intersubjectively stable, publically accessible tree that can be identified and re-identified over a sufficiently long run and whose identity conditions supervene on the sensory and cognitive states of the agents inhabiting the associated VE.

When typical inhabitants of a MMORPG confront a tree, and say, "There is a tree," what they are talking about is best captured by 7.

The ontological status of virtual objects

I have just argued that the essential nature of virtual trees is best described not in terms of the mind-independent states of physical systems, but in terms that make essential reference to the cognitive states of human agents. One might reasonably use either the language of a verificationist semantics as embraced by some antirealists, or the language of conscious, first-person phenomenal states familiar from recent defenders of phenomenal consciousness. I have purposely used both in this discussion, not wanting to privilege either one. Some will likely find one option more congenial than the other.

I recognize that I have done nothing to answer the myriad objections that can quite reasonably be raised against this controversial position. But that work must be left to another time. Our final task now is to come full circle on the Hacking quote that opened this paper. If (by chance) I am right about the essential nature of virtual objects, are they real or not? Do we add them to our ontology, or leave them out?

Lynn Rudder Baker has, in the pages of this publication (Spring 2008¹⁰), addressed a question that is, at least in part, relevant our question: Are artifacts less real than natural objects because they are mind-dependent? In that article, she defends a position with which I am completely sympathetic. She insists that human artifacts are in no way metaphysically deficient in virtue of their having been shaped and fashioned by the human mind. Tables are no less metaphysically "real" than are tubers. She says,

There is a venerable—but, I think, theoretically misguided—distinction in philosophy between what is mind-independent and what is mind-dependent. Anything that depends on our conventions, practices, or language is mind-dependent (and consequently downgraded by those who rest metaphysics on a mind-independence/mind-dependence distinction)...

A second reason that the mind-independent/mind-dependent distinction is unhelpful is that advances in technology have blurred the difference between natural objects and artifacts. For example, so-called "digital organisms" are computer programs that (like biological organisms) can mutate, reproduce, and compete with one another. ...Are these objects... artifacts or natural objects? Does it matter? (p.4)

I wholeheartedly affirm Rudder Baker's sentiments here. A chair is not dependent upon the *present* cognitive activity of any agent. Every mind that exists in the universe could cease to exist and the chair would continue to exist. The chair is *causally* dependent upon the *past* activity of some cognitive agent, but so long as that past activity produced something with its own intrinsic properties, which is dependent upon no *present* mental activity, then it is mind-independent in the metaphysically relevant sense.

Those who Rudder Baker criticizes make the mistake of conflating two fundamentally different meanings of the term "mind-dependence." The kind of mind-dependence attributable to artifacts which are dependent only in their *causal origins* has profoundly different ontological implications than the kind of mind-dependence attributable to virtual objects which will literally *cease to exist* if minds stop thinking about them. Treating the former as if it deserves the same ontological classification as the latter ignores this important distinction. That is why, according to my definitions, artifacts qualify as physical objects (PO) not virtual objects (VO).

Having said that, I would argue that Rudder Baker makes too sweeping a claim in her concluding two sentences of the paper, where she says:

No one who takes artifacts of any sort seriously, ontologically speaking, should suppose that metaphysics can be based on a distinction between mind-independence and mind-dependence. In any case, technology will continue to shrink the distinction, and with it, the distinction between artifacts and natural objects. (*Ibid.*)

Ironically, Rudder Baker seems to be committing the same conflation error as did her opponents, but in the reverse direction. She seems to be denying that there could be any ontologically significant distinction between mind-dependence and mind-independence. The table in my kitchen, and the table in my *Second Life* kitchen are both artifacts. But one is mind-dependent in a more robust sense that *is* ontologically significant. If she is denying this, then I think she makes a mistake similar to her opponents. In the end, I am not at all confident that she is denying this distinction. She may simply be ignoring it. But if she is denying it, then we deserve more of an argument than she has given thus far.

Conclusion

So what about virtual trees? Are they real? The term "real" is ambiguous. No, they are not real in the sense that is opposed to being ideal, or verificationist. They are metaphysically dependent upon the cognitive state of human beings, and that makes them "ideal" in contrast to physical trees that are "real."

But more important than being “real,” in this sense, is being real in the sense that they “exist” in some substantive sense of that word. Yes, they do exist. They are real enough to make it the case that statements like “I left your shield next to the tree in the quad” and “You lied to me about the power of this sword and I will never be your friend again” are *true*. These statements are not about a fictional realm like that of a novel. One can commit real (not fictional) betrayal with a virtual sword, and a sword real enough to betray is real indeed.

I also believe it is reasonable to say that virtual trees exist and should be listed among our ontological commitments in something like the way that *conscious states* should be listed in one’s ontology.¹¹ The conscious states of John’s believing that *p* and seeing an orange sunset are real. So too, the virtual tree that supervenes on the conscious states of many people. Obviously, those who deny the existence of consciousness will not be persuaded by this line of reasoning, but for that there is also the verificationist route to metaphysical legitimacy.

This discussion has only begun to explore the ontological status of virtual objects. It is a discussion that I hope continues.

Endnotes

1. May 1, 2006, *Business Week* cover title “Virtual World, Real Money.”
2. To learn more about the very real business of buying, selling, and owning virtual objects, see the online newsletter, *Virtual Goods News*, <http://www.virtualgoodsnews.com>.
3. Ian Hacking. *Representing and Intervening* (Cambridge: Cambridge University Press, 1983), 23.
4. I don’t mean to engage any philosophical disputes surrounding the nature of “space-time,” if only because I have no competence in that domain. I am hoping for a conception as neutral as is possible, believing that nothing of significance hangs on it.
5. Michael Dummett. *Truth and Other Enigmas* (Cambridge, MA: Harvard University Press, 1978).
6. Hilary Putnam. *Reason, Truth & History* (Cambridge: Cambridge University Press, 1983), 1-21.
7. For more on the arguments of Putnam and Dummett and their relationship to the realism-antirealism debate, see David L. Anderson, “What is Realist about Putnam’s Internal Realism,” *Philosophical Topics* 20 (1992): 49-84.
8. *Ibid.*
9. Those who embrace the Identity thesis for the nature of mental states might want to argue that condition 6. more accurately describes *the very same reality* as condition 7. Since I reject the Identity thesis, I don’t hold this view but in the present context I wouldn’t feel compelled to argue with such a person. Condition 6. is close enough to victory for me.
10. Lynn Rudder Baker. “The Shrinking Difference between Artifacts and Natural Objects.” *American Philosophical Association Newsletter on Philosophy and Computers* 07 (Spring 2008): 2-5.
11. For more on the reality of conscious, see David Leech Anderson, “Consciousness & Realism” *Journal of Consciousness Studies* 14 (2007): 1-17.

Realism and Antirealism in Informatics Ontologies

Robert Arp

National Center for Biomedical Ontology, University at Buffalo

Abstract

The realism-antirealism debate in the philosophy of science has made its way into informatics and computer science circles in debates concerning the status of the entities represented in what informaticians call ontologies. For realists, the terms of these ontologies refer to real entities out there in the world; for antirealists, they refer to concepts in the minds of experts. In this paper, after an explanation of domain and formal ontologies, I offer some criticism of the antirealist approach and argue that, in spite of the antirealist sentiments that still predominate in informatics circles, informaticians can nonetheless feel comfortable in constructing domain and formal ontologies from a realist perspective.

Key Words: informatics, ontology, domain ontology, formal ontology, realism, antirealism

The Sea of Information

Informatics is the science associated with the collection, categorization, management, storage, processing, retrieval, and dissemination of data and information—principally, through the use of computers as well as computational and mathematical models—with the overall goal of improving retrieval and dissemination of data and information. Increasingly, many more traditional disciplines have their own informatics, reflecting the fact that they are confronted by the need to deal with large bodies of data and information—consider, for example, the field of Geographic Information Systems (<http://www.gis.com/>) or of biomedical informatics (Shortliffe & Cimino 2006).

The body of information deriving from such disciplines that is now being made freely available through computers on the Web constitutes a veritable sea of extraordinary depth and breadth. How can we collect, categorize, manage, store, process, retrieve, disseminate, mine, and query all of this data and information appropriately and efficiently by computational means?

The problem is to chart this ever-growing sea of information in such a way that its various parts can be efficiently accessed, used, navigated, and reasoned about by human individuals. How can we ensure that the terminology, definitions, relations that are used when storing information and data (a) accurately reflect the developing state of knowledge in a particular domain or discipline, (b) are internally coherent, (c) are clearly defined, and (d) are interoperable from one database to the next?

Here, it is especially (d) that poses problems. Researchers in different disciplines speak different languages, use different terminologies, and format the results of their research in different ways. The situation is not unlike that of the Biblical Tower of Babel, where there is an uncontrolled and unsurvivable multiplicity of different languages and little in the way of cross-linguistic understanding. Because bodies of data are insulated from each other in this way, interoperability, shareability, and reusability of data and information is greatly limited. The result is a *silo effect*: data and information are isolated in multiple, incompatible silos. And it is to address the silo problem, philosophers, computer scientists, and informaticians in various domains have worked to create what are known as *domain ontologies* in their respective fields of study.