Proceedings of the European Society for Aesthetics

Volume 6, 2014

Edited by Fabian Dorsch and Dan-Eugen Ratiu

Published by the European Society for Aesthetics

esa

Proceedings of the European Society of Aesthetics

Founded in 2009 by Fabian Dorsch

Internet: http://proceedings.eurosa.org
Email: proceedings@eurosa.org
ISSN: 1664 – 5278

Publisher
The European Society for Aesthetics

Department of Philosophy
University of Fribourg
Avenue de l'Europe 20
1700 Fribourg
Switzerland

Internet: http://www.eurosa.org
Email: secretary@eurosa.org

Proceedings of the European Society for Aesthetics

Volume 6, 2014

Edited by Fabian Dorsch and Dan-Eugen Ratiu

Table of Contents

The 'Uncanny Valley' and Spectating Animated Objects

James R. Hamilton*
Kansas State University

Abstract. The thing that strikes most thinkers about puppets and other animated objects in theatre is that they can seem genuinely 'uncanny'. So, instead of thinking first about how spectators grasp what is going on in any performance and fitting the exceptional experience of the uncanny within that larger story, most theorists I know about have begun first with an account of the uncanny nature of puppets—or of something that would readily explain it—and only then worked out a general theory of spectating puppets and animated objects based on their preferred account of the exception. Although it might seem natural to begin with what is most striking, given the ubiquity and prominence of that uncanny feeling, in my view that approach is exactly backwards. In this paper I show why that is so and suggest some lines of research that might help to redress the situation.

1. *Introduction*

When, if ever, do animated objects trigger the 'uncanny valley' effect? This effect is thought to be a combination, perhaps a blending, of repulsion and attraction, an 'unsettling delight' (Gross, 2011: 2) felt in the presence of figures whose visually apparent features are very close to, but not exactly like, those of a healthy human being. When robotics engineer and designer Masahiro Mori first described this effect (hereinafter the 'UV effect'), he initially conceived of it as a challenge in robotics design (Mori, 1970/2012). Since then, and very quickly, the notion took hold in the field of visual animation as an explanation for the success and failure of various animated feature films and the appeal or repulsion of figures in video games. And the UV effect is now a standard accepted element in the design vocabulary of animation and robotics.

* Email: hamilton@ksu.edu

One might expect this concept also to play a role in describing some aesthetic properties of puppets and other animated objects. For one thing, it is already a commonplace among those who write on puppets to describe them as 'uncanny objects' (Gross, 2011; Simms, 1996; Zamir, 2010). Many of the deeply interesting aesthetic properties of animated objects—such as those arising from their very materiality and those affecting spectator responses even when the objects are not involved in a performance—seem to have something to do with their sheer appearance and spectator reactions to it. So, a *reasonable hope* is that many of those deeply interesting aesthetic properties of puppets and other animated objects will receive an explanation as a kind of 'spin-off' from the explanations that can be given of the UV effect in the presence of animated objects.

In the first section of the paper, I describe the UV effect in greater detail. I also present a rough account of the state of play, so to speak, in the robotics and animation literature on what may be a range of phenomena.

In the second section of this paper, I examine a class of views that are promising candidates for making good on the *reasonable hope* expressed above. These views (represented here by Zamir, 2010) rest on a familiar metaphysical distinction between objecthood and subjectivity, holding that puppets just *are known to be* objects but that they *seem to be* subjects. The resulting epistemic disparity, between our knowledge of the object status of puppets and our experience of them as subjects, gives rise, on these views, to the recognition of some uncomfortable facts about our interactions, attempts to control, and desires towards (even the desire to *be*) matter. Moreover, the views are able simultaneously to connect with the relevant aesthetic issues and to offer an account of the UV effect. But, as I show at the end of the section, this class of theories claims too much, in one sense, and in another too little, insofar as it simply fails to be generalizable beyond puppets to other animated objects.

In the third section I argue that theories belonging to this class also obscure both the distinction and the epistemic relations between data and inferences spectators make on the basis of data. And I assemble some empirical evidence that shows these are clearly distinct from one another in important ways. However, I also argue that we do learn from considering these views because they are able to show us why any alternative explanation should begin by giving a general account of spectating pup-

petry, should connect with important aesthetic issues, should meet the plausible requirement of generalizability, and should make plain the distinction and the relations between data and inference. These form a set of *desiderata* for any alternative theory. Moreover, as I argue in this section, any alternative proposal should be consistent with a broad set of results at the intersection of empirical cognitive science and formal learning theory; and this is a final *desideratum*.

In the final section of the paper, I sketch the outlines of an alternative approach that would meet all the *desiderata*. I argue that the strategy lying behind the view must allow us to re-describe the UV effect as a specific class of cognitive effects having fairly immediate affective consequences and a specific kind of etiology.

But first, one short note on terminology. Following Cariad Astles (2009) and others (see Furse, 2008), I will write about puppets, generally, as instances of 'animated objects'. Ever since Frank Proschan's groundbreaking article (1983), some people customarily refer to these as '*performing* objects' (for example Bell, 1997; and Cohen, 2007). The referential scope of the term 'performing objects' may be broader than that of 'animated objects'. For it can be used to refer to elements in the *mise en scene* that are not even animated but are nevertheless important to the experience of spectators and, so, figure importantly in their cognitive and affective uptake of a performance. I have no reason, other than perspicuity regarding the issues discussed herein, for choosing to use the term 'animated objects' in its stead. So, if you prefer the other term, feel free to substitute it in. Nothing I argue for in this paper hangs on whatever other differences might follow from classifying these objects under those different category terms. For the most part, I will use the term 'animated objects', primarily because the issues I address about puppets and other *animated* performing objects are made somewhat clearer when described using it.

2. *The UV Effect*

Let us begin with one popular way to describe the Uncanny Valley effect.

> [...] Freud popularized the idea of the uncanny, the blend of attraction and repulsion we feel for something we can't quite categorize. A Japanese engineer adapted the notion in the 1970s for work in robotics, and the idea was later extended to animation. In short, it says that if artists created characters that were only vaguely human, like cars with faces or anthropomorphic ducks, viewers found them endearing. Similarly, if artists drew realistic characters of almost photographic quality, viewers also adored them. Something funny happened between the two extremes, though. If a drawing or robot looked mostly human but not quite, it actually repelled people. Computer-generated characters in movies often tumble into this uncanny valley, not to mention zombies, clowns and celebrities with bad face-lifts. It seems that when something is, say, 50 percent human, our brains focus on the similarities and we embrace it. When it's 95 percent human, we focus on the differences, and the unresolved conflict we feel — is that human or not-human? — creeps us out. (Sam Kean, 2014)

The Japanese engineer referred to in this quotation was Masahiro Mori who is a major figure in robotics design. It was precisely because of this very reaction, the 'creepiness' Mori felt when working with some robots, that he counseled robotics designers to avoid creating robots that would fall into the uncanny valley. Mori believed robotics engineers, who strive to make robots with which human beings can effectively interact, could achieve their goals without trying to make robots look like human beings. And, although this recommendation does have its doubters (Hanson, *et al*., 2005), this strategy of avoiding the uncanny valley still dominates in both robotics design and animation, including the design of animated video games.

One important aspect of the Mori's account of the UV effect that gets obscured by Kean is that, as Mori had more clearly emphasized, the UV effect is something observers experience on the basis not only of physical appearance but also on the basis of *movement*. That is, Mori held, the UV effect is actually heightened, made more intense, if the object perceived

as uncanny also moves. This is evident in the graph developed to translate the one first presented by Mori.

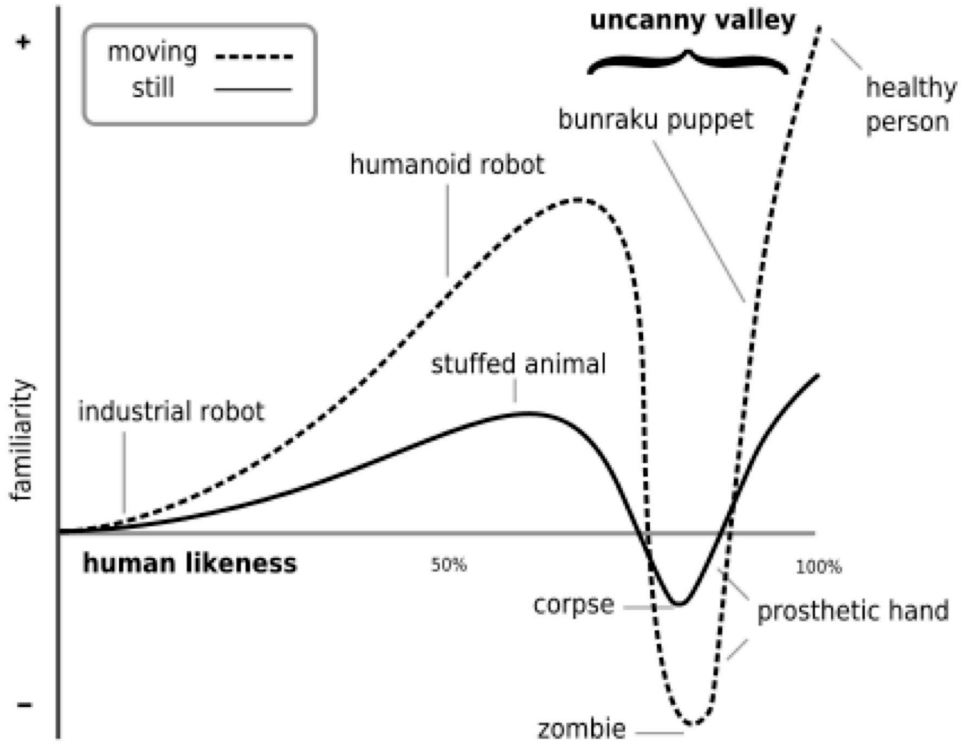

Figure 1. Source: M. Mori, 2012.

On Mori's hypothesis [Figure 1], as the appearance of an object approaches that of a healthy person (the 'human likeness' indicated on the horizontal axis), the familiarity an observer experiences is increasingly positive (vertical axis), that is, until a sudden drop off in familiarity is reached when the experience becomes strikingly negative (hence the term 'valley'). The dotted line represents movement. And the contrast between the dotted and solid lines shows that the experience of familiarity that an observer has is enhanced by perception of motion, both positively and negatively.

But this nice, popular, and deeply intuitive hypothesis is in serious trouble.

Mori rested the hypothesis only on anecdotal evidence and never tested it further. In fact, it was not empirically tested until some twenty years later. And when it was, a number of questions emerged. Not least was the question whether there is a genuine phenomenon here at all (Hanson, *et al*., 2005). Other questions soon followed: Is it a *single* phenomenon? Is it triggered *only* by animated objects, human-like robots, and/or video animations? Is it triggered *only* at the 95-98% part of the curve? And none of these interesting questions has yet, so far as I know, received a convincing answer (Guizzo, 2010).

However, rather than pursue these skeptical questions here, what I want to do is assume the original story is roughly correct and then look for potential causes of the phenomenon. In the end, I believe, this strategy pays off by showing us what needs to be done first in accounting for the relevant reactions we have to puppets and other animated objects.

3. *A Promising Class of Theories*

In this brief section of this paper, I outline and discuss a class of theories that have seemed to be promising candidates for what we need. This class is represented here by Tzachi Zamir's essay, simply entitled 'Puppets'. As a class they rest on a familiar metaphysical distinction between objecthood and subjectivity, namely that puppets are objects that seem to be subjects. The resulting epistemic disparity gives rise to or engages a tension that results in the common ascription of uncanniness to puppets. A reasonable hope lying behind the class of views is that many interesting aesthetic properties of puppets might receive an explanation as a kind of 'spin-off' from the explanations that can be given of the UC effect and the experiences it subsequently engenders.

What are those experiences? Zamir encapsulates them this way: the phenomenon—the epistemic disparity between our knowledge of the object status of puppets and our experience of them as subjects—forces us to recognize a set of uncomfortable facts: about the 'uncontrolled and uncontrollable nature of matter'; about subjectivity as 'life *qua* momentarily resuscitated matter'; of 'the illusiveness of freedom and the disturbing autonomy possessed by [our] creations'; and of the fact that *we* sometimes

want to become objects, *and that we like it* (Zamir, 2010: 389, 392, 392, and 394-395). So, then, here is how this goes: as a direct result of recognizing these facts, this class of views is able to connect with and appears to provide answers to some important aesthetic questions.

By 'aesthetic questions' I mean questions of two kinds. First, aesthetic questions concern how to characterize factors that generate what can and typically do figure into descriptions, interpretations, and evaluations of uncontroversially recognizable aesthetic experiences or the objects that cause them. Second, aesthetic questions concern how, precisely, these factors *do* figure into such descriptions, interpretations, and evaluations. Such experiences are frequently, but not necessarily, generated by features of works of art. But they are frequently generated by features of naturally occurring objects and events as well. Moreover, nothing in this idea presupposes that the experiences themselves are to be regarded positively. One can, for example, experience horror and like it; and one can just as easily experience feelings of warmth and security but dislike it.

Among the aesthetic questions generated by the experience of puppets are questions about how to characterize and explain the following factors: (a) that makers, users, and spectators are fascinated with the *materiality* of animated objects; (b) that many spectators experience strong emotional reactions to animated objects when they are *not* part of a performance (fiction-making or otherwise); and (c) that puppets seem to have a life of their own, both *before,* on-stage, *and afterwards* when they appear (as they sometimes do) in dressing rooms, museums and galleries.

To reiterate, it would seem that a view of our experience of puppets that is grounded in the tension we might feel—a result of the epistemic disparity already described—would be able to deliver a plausible set of answers to each of these aesthetic questions. Thus, that must be taken seriously as a real advantage of that approach to these aesthetic questions.

Unfortunately, it appears the approach, the entire class of such theories, is too strong and fails to be generalizable. First, while it seems able to deliver a story about the UC effect and puppets and, based upon that story, also appears able to tell how certain kinds of experience are caused as well as what those kinds of experiences are likely to predict by way of answers to some important aesthetic questions, the price paid by the approach is too high. For the story it has to tell about puppets is simply that

they just *are* uncanny objects, *all of them are,* and *all the time*. But that is not consistent with the facts. Not all puppets are regarded as uncanny, nor are they so regarded all of the time. We will come back to this point.

Secondly, this class of theories fails to be generalizable beyond puppets to other animated objects where animacy, but *not* subjectivity, is what is expressed or experienced. Precisely because it rests on the metaphysical distinction between objects and subjects it cannot tell us much about those items—increasingly used in theatre—that are not subjects but only animated. Walls that breathe need not be subjects, nor experienced as such in order to produce interesting, compelling, and yes even 'uncanny' theatrical experiences in spectators. Nor is it clear what that they have to be understood as *expressing* anything so much as they are taken to be generating *moods*. And moods form a class of feelings that are much more diffuse than ordinary emotions, which of course *can be* expressed. Moreover, not all acts of creation or conveyance—as in, 'the company created and conveyed a mood of melancholia'—are also acts of expression, *per se*.

4. *A General Plan to Recover*

Crucially, I now argue, theories in the class we have been considering also obscure both the distinction and the epistemic relations between data and inferences spectators make on the basis of data. Moreover, as I will also argue later in this section, any alternative class of theories or hypotheses that explains the UV effect should be consistent with a broad set of results at the intersection of empirical cognitive science and formal learning theory.

To see what I have in mind by the first contention, consider some empirically derivable distinctions. Consider first the distinction between detecting causation and detecting animacy. This was first studied systematically by Heider and Simmel in 1944. This was followed up by more detailed work by Albert Michotte in 1946. Michotte's tests were simple. Michotte devised a way to manipulate the movement of two objects on a projection screen. Both could move left and right at various speeds and with various delays. In the 'launching experiments', for example, a ball on the left moved from left to right about 20 cm, then stopped for about 3

seconds; and then the ball on the right moved about 10 cm from left to right and stopped. In 'rest to motion' experiments, a single ball would be seen at rest and then move for a certain distance.

Although the tests themselves were simple, the results of the tests were extremely powerful and sophisticated. One measure of this was that very small changes in either the distance travelled or the timing of pauses reliably produced substantial variations in the descriptions of what subjects saw. In particular, whereas the launching, entraining, and expulsion examples prompted most subjects to see some object(s) causing other objects to move, the tests involving subjects responses to rest to motion, apparent goal-directedness, and apparent collision avoidance strongly suggest most subjects see some objects moving in some patterns as self-movers, as (at least) *animated*. The now classic treatment of these issues and the tests that are now most often cited is given in Castelli, Happé, Frith, and Frith (2000). And what, in the present context, is most remarkable about that treatment is that all the tests involve animated objects in the form of the mere representation of abstract objects such as variously colored triangles, squares, rectangles, and circles, reproduced on video screens (Scholl and Tremoulet, 2000).

A second empirically derived distinction is between perception of causation and causal inferences. The terms used for these in the relevant research literature are 'perceptual causality' and 'reasoning to underlying mechanisms'. The rough idea is that while even very young infants detect causation among events perceptually, it takes the development of reasoning skills before a child is able to reliably predict, for example, when a pair of events that look like they involve causation do not in fact do so (Schlottman and Shanks, 1992); Schlottman and Surian, 1999; Schlottman, 1999: and Castelli, Happé, Frith, and Frith, 2000).

A third empirically derivable distinction is that among attributions of causation, attributions of animacy, and attributions of beliefs. The latter is often referred to now as the attribution of a 'Theory of Mind'; and it is arguably the case that only when a 'ToM', or something very like it, is attributable that we have genuine—even if only minimal—'subjectivity'. The research on this began in the early 1980s, and is now part of the standard canon of empirical social science (Wimmer & Perner, 1983; Baron-Cohen, Leslie, & Frith, 1985; Johnson, Slaughter, & Carey, 2002; Lohman,

Carpenter, & Call, 2005; and Hedger & Fabricius, 2011). Two features of the experiments purporting to support the existence of ToM in children beyond infancy are especially worth noting in the present context. The first is that the initial tests to determine when children will reliably predict the behavior of others by attributing beliefs to them—the 'false belief' test—were conducted using puppets (famously called 'Sally' and 'Ann'). The second is the important finding, in the work of Baron-Cohen and his colleagues, that failure to pass the false-belief tests is characteristic of children with autism. Indeed, much of the subsequent work on the subtleties of ToM has been conducted with this very practical orientation in mind.

These studies from the current psychological literature show us three things. First, only certain patterns of movement are responsible for engendering the perceptions of causation, of animacy, and of subjectivity, respectively. This entails there are important and describable differences among them. Second, the patterns of movement responsible for engendering those different responses, both in perception and in reasoning, can be produced using clearly *in*animate objects that can be made to behave as though animated or as though subjects. And, finally, these references also should convince us that, while some of the time the recognition of causation and of animacy is perceptual, the recognition of subjectivity is *never* purely perceptual and always involves taking very seriously the relevance of *both* the kinds of data a person is presented *and* the person's capacity to draw inferences from that data.

What, then, have we seen so far? Some of the promise of the class of theories we considered in the previous section consisted in its ability to connect with important aesthetic issues. Some of it consisted in its apparent ability to give a simple and direct account of the UV effect. But it failed to deliver a general account of spectating puppetry because it claimed too much and failed a plausible requirement of generalizability. In this section we have seen is what it obscures, namely, both the distinction and the epistemic relations between data and inference. So, in the end, it appears to wrap a puzzle in a mystery and call *that* an explanation.

However, studying that class of theories also reveals something positive. In particular, what we learn about explaining the UV effect is, I believe, that explaining the UV effect itself should probably not be our starting point but, rather, one of the elements to be explained by a more

general theory of the nature of spectating animated objects. That is, we learn from examining that class of theories that any alternative explanation should first give a general account of spectating puppetry, should connect with important aesthetic issues, should meet the plausible requirement of generalizability, and should make plain the distinction and the relations between data and inference.

Moreover, I now argue, any proposal alternative to the metaphysically grounded class of theories presented and discussed in the previous section should *be consistent with* a broad set of results at the intersection of empirical cognitive science and formal learning theory.

Why is that? The intuition here rests on the simple idea that explanations track causes. There are several accounts of what makes a body of statements *explanatory* of some phenomena. Most people who think about explanations, *per se*, agree that—for physics, chemistry, biology, psychology, anthropology, and so on—there has to be some causation or at least something like causal relevance involved in order to have an explanation. From that basis of agreement, there is significant divergence (Halpern and Pearl, 2005a and 2005b, and Christopher Hitchcock, 2007).

Some philosophers of science think that, ultimately, the structure of an explanation should be deductive, where the laws, theorems, and so on, function as premises in the argument, some general observed facts also figure as premises, and the conclusion is the phenomenon to be explained. The phenomenon in question may have occurred and stand in need of explanation. Or it may not have occurred yet, in which case it is a prediction based upon that explanation. Indeed, this idea suggests a method of testing an explanation, namely, that one should first see what the explanation, taken together with known facts, allows one to predict and then second run experiments to see if the predicted phenomena actually occur. These philosophers allow that some explanations are of sufficiently particular phenomena that the deductive model does not work for them; in such cases they think we use inductive arguments that take laws (and the like) as premises and offer something like statistical syllogisms to arrive at the phenomenon to be explained.

Others philosophers of science have noted the historical fact that scientific explanations have almost always been inductive. They involve induction over sets of data (samples of phenomena that exist and can be

given clear descriptions) that leads to something like generalizations about the whole set of phenomena (a whole population). The question is how to capture certain features of a situation in terms of the notion of statistical relevance or conditional dependence relationships. For an explanation, remember, cannot simply be a correlation among data sets, it must involve causal relations or at least causal relevance among the features.

Still other philosophers of science, and sociologists of science in particular, take the issue to be one that is fundamentally social in nature. What an explanation is, on one of these influential views, is a unified account of a range of different phenomena that yields a sense of understanding among the explaining community. This is clearly an intuitively appealing idea, since it trades on the idea that a good theory offers accounts of phenomena we have not have anticipated belonging to the theory. As such, it offers to unify phenomena in unanticipated ways. Historically, again, this fact about theories has played an important part in the development of scientific theories. It is unclear how well it applies to theories outside of the sciences, however.

Perhaps, of course, there are different kinds of explanations, suitable for different ranges of phenomena. This would imply, of course, that there are different kinds of theories as well. And determining how to specify those is one of the unsolved problems in philosophy. But, to reiterate the main point, whatever else goes into the explanatory nature of a theory, the theory had best track causes or it won't pass muster at all.

This is why it seems to me that the resources we really need for understanding how spectators appreciate animated objects, and how we are to explain the UV effect as well, are best found not in a metaphysical distinction, however intuitively persuasive it otherwise might be, but in empirical psychological research and formal learning theory about how we learn to track causes.

5. *An Alternative Strategy*

In this final section of the paper, I quickly present and partially explain an alternative strategy for developing a view that meets the foregoing *desiderata*. And I argue that any such strategy should allow us to re-describe

the UV effect as a specific class of cognitive effects having fairly immediate affective consequences and one or more specific kinds of etiology.

a. Some of the Relevant Empirical Research

As a first step we think of the general term 'experience' more precisely as a kind of perceptual input. But, for various reasons, most to do with the bewildering variety of perceptual inputs an individual experiences at any given time—not to mention the difficulty in moving from such pure 'raw' data to anything comprehensible—we specify the kind of perceptual input as that which is the result of 'perceptual input analyzers' (Carey, 2009).

What does this mean? As Carey conceives them, perceptual input analyzers are modules that perform simple transformations of raw perceptual input, make raw perceptual input suitable for conceptual cognition, and make latent abstract concepts more salient and more learnable. A final point, perceptual input analyzers may themselves be learned (Saxe and Carey, 2006).

Crucial to this approach is the fact that perceptual input analyzers place one kind of constraint on the sheer size of hypothesis spaces an agent must be thought of as considering when confronted with the world.

Because it has this effect and because of the way it does this, one might well be reminded of a similar strategy using the concept of 'affordances'—a concept owing its initial provenance to Jerome Gibson (1979) and since developed by Shaun Gallagher, Daniel Hutto and others. Of course, some dissident views have focused especially on two claims Gallagher and Hutto have made: that the approach represents a 'second generation of cognitive science'; and that the approach demonstrates the inadequacy of ToM (Lawrence Shapiro, 2007; Shannon Spaulding, 2010 and 2011). But for the present purposes, it does not matter which side one takes on those finer-grained issues, so long as we conceive of perception as, in some sense but at a very basic level, 'ready for conceptualization' that place constraints on the sheer size of hypothesis spaces an agent must entertain.

b. The Shape of the Relevant Formal Learning Theory

On the approach I recommend, spectators experience 'pre-analyzed' data and infer causal theories to explain them and/or to discover categories

within which the data fall. The metaphysical approach considered in §2 focuses only on the latter kinds of inference, but as we have seen, does not actually treat that as an *inference* at all. This is part of why it obscures the relationship between data and inferences.

But it altogether ignores the inference to the fact that agents attempt to discern which causes (and which kinds of causal theories) could explain the data stream. In contrast, on the approach I recommend, we do assume a distinction between inferences to particular causes (the inference of causal theories) and inference to causal or category structures. Both are in the form of considerations concerning hypotheses one might entertain in order to explain one's experience. But they are different kinds of hypotheses. This is beneficial because most agents are adept at quickly inferring particular causal theories/categories from data streams.

What is now and only recently realized, at least in formal modeling theory, is that causal and categorial *structures* are inferred from the same data as particular causes are and at nearly the same time. Moreover, these facts are connected (Tenenbaum, *Et al.*, 2011 and Goodman, *et al.,* 2011).

c. The Alternative Strategy

So, what should an alternative strategy aim at doing? First, it would redescribe the UV effect as a cognitive effect that is one thing that can happen when a spectator is trying to understand a performance, whether by human actors, puppets, or other animated objects. The kind of thing that can happen will occur as a problem with inferring the causes that generate the stream of our experiences or as a problem with inferring the category of objects to which we take an animated object to belong. Although much of the time, a spectator will draw defensible and plausible inferences about what is happening and what kinds of things she is confronted by, this comparatively 'normal' process can be interrupted and disturbed. If so, problems arise in the inferences to appropriate hypotheses. The UV effect, if it ever occurs, would originate for some particular spectators as a problem of this kind.

Second, we would expect an investigation that takes this approach to seek to determine which further affective consequences might attend to or be generated by these problems in hypothesis formation. We already

know, from the experiences of puppet-makers, puppeteers and those who make and use other kinds of animated objects in theatre performances that the relevant affective consequencs are probably fairly immediate and we now know they must express a spectator's confusion about the *kinds* of causes/categories that might make sense of our experiences.

Thirdly, we might expect, and should certainly investigate, whether there is a specific etiology that yields this kind of confusion. That is, as many have suggested, the ultimate causes must be understood as stemming directly from the experiencc of encountering animated objects or from encountering human beings in unexpected situations or possessing certain unexpected features. They are not, or not primarily, due to our having some sort of psychological set, bias, or propensity to draw mistaken inferences before we encounter an animated object. Nor are all animated objects going to produce the effect, at least not all the time. They are rather caused by our encounter with those precise objects, or better, with precisely specifiable features of those objects. In the approach I recommend, that would be theorized as caused by a lack of clear presentation of unexpected features, of certain specific sort(s), in the data stream itself.

References

Astles, Cariad (2009), 'Wood and Waterfall: Puppetry Training and its Anthropology', *Performance Research,* vol. 14 (2), pp. 54-59.

Baron-Cohen, Simon, *et alia* (1985), 'Does the autistic child have a "theory of mind"?' *Cognition,* vol. 21, pp. 37-46.

Bell, John (1997), 'Puppets and Performing Objects in the Twentieth Century', *Performing Arts Journal,* vol. 19 (2), pp. 29-46.

Carey, Susan (2009), *The Origin of Concepts*: *Oxford Series in Cognitive Development*, Oxford: Oxford University Press.

Castelli, Fulvia *et alia* (2000), 'Movement and Mind: A Functional Imaging Study of Perception and Interpretation of Complex Intentional Movement Patterns', *NeuroImage,* vol. 12, pp. 314-325.

Cohen, Matthew Issac (2007), 'Puppetry and the destruction of the object', *Performance Research,* vol. 12 (4), pp. 123-131.

Furse, Anna (2008), 'Committed to the Other: The ethical nature of work with puppets', *Puppet Notebook,* vol. 13: 19-21.

Gallagher, Shaun (2001), 'The practice of mind: theory, simulation, or interaction?' *Journal of Consciousness Studies*, vol. 8, pp. 83–107.

— (2004), 'Understanding problems in autism: interaction theory as an alternative to theory of mind', *Philosophy, Psychiatry, and Psychology*, vol. 11, pp. 199–217.

Gibson, Jerome J. (1979), *The ecological approach to visual perception*, New York: Houghton Mifflin.

Goodman, Noah , *et alia* (2011) 'Learning a theory of causality', *Psychological Review*, vol. 118 (1), pp. 110-119.

Gross, Kenneth (2011), *Puppet: An Essay on Uncanny Life*, Chicago: University of Chicago Press.

Guizzo, Erico (2010), 'Who's Afraid of the Uncanny Valley?' *IEEE Spectrum* (on-line), accessed at: `http://spectrum.ieee.org/automaton/robotics/humanoids/040210-who-is-afraid-of-the-uncanny-valley`, on June 19, 2014.

Halpern, Joseph Y. and Pearl, Judea (2005a), 'Causes and Explanations: A Structural Model Approach; Part I: Causes', *British Journal for the Philosophy of Science*, vol. 56, pp. 843–887.

— (2005b), 'Causes and Explanations: A Structural Model Approach; Part II: Explanations', *British Journal for the Philosophy of Science*, vol. 56, pp. 889–911.

Hanson, David, *et alia* (2005), 'Upending the uncanny valley', *20th National Conference on Artificial Intelligence*, vol. 4, pp. 1728–1729.

Hedger, Joseph A. and Fabricius, William V. (2011), *The Review of Philosophy and Psychology*, vol. 2 (3), pp. 429-447.

Heider, Fritz and Simmel, Marianne (1944), 'An Experimental Study of Apparent Behavior', *The American Journal of Psychology*, vol. 57 (2), pp. 243-255.

Hitchcock, Christopher (2007), 'Three Concepts of Causation', *Philosophy*

Compass vol. 2 (3), pp. 508–516.

Hutto, Daniel D. (2004), 'The limits of spectatorial folk psychology', *Mind & Language,* vol. 19, pp. 548–573.

Johnson, Susan, Slaughter, Virginia and Carey, Susan (2002), 'Whose gaze will infants follow? The elicitation of gaze following in 12-month-olds', *Developmental Science*, vol. 1 (2), pp. 233-238.

Kean, Sam (2014), 'Beyond the damaged brain', *New York Times*. Accessed on May 3, 2014 at `http://nyti.ms/1iTxL0t`.

Lohmann, Heidemarie, Carpenter, Malinda and Call, Josep (2005), 'Guessing versus choosing – and seeing versus believing – in false belief tasks', *British Journal of Developmental Psychology*, vol. 23, pp. 451–469.

Michotte, Albert (1962), *The perception of causality*, Andover, MA: Methuen.

Mori, Masahiro (1970), 'The Uncanny Valley,' *Energy*, vol. 7 (4), pp. 33–35 (in Japanese) and Mori, Masahiro (2012), 'The Uncanny Valley,' translated by MacDorman, K. F. and Kageki, N., *IEEE Spectrum* (on-line), accessed at: `http://spectrum.ieee.org/automaton/robotics/ humanoids/ the-uncanny-valley`, on June 7, 2014.

Proschan, Frank (1983), 'The semiotic study of puppets, masks, and performing objects', *Semiotica*, vol. 1 (4), pp. 3-44.

Saxe, Rebecca and Carey, Susan (2006), 'The Perception of Causality in Infancy', *Acta Psychologica*, vol. 123, pp. 144–165.

Schlottmann, Anne (1999), 'Seeing It Happen and Knowing How it Works: How Children Understand the Relation Between Perceptual Causality and Underlying Mechanism', *Developmental Psychology*, vol. 35 (5), pp. 303-317.

Schlottmann, Anne and Shanks, David R. (1992), 'Evidence for a Distinction between Judged and Perceived Causality', *The Quarterly Journal of Experimental Psychology*, vol. 44(A) (2), pp. 321-342.

Schlottmann, Anne and Surian, Luca (1999), 'Do 9-month-olds perceive causation-at-a-distance?' *Perception*, vol. 28, pp. 1105-1113.

Shapiro, Lawrence (2007), 'The Embodied Cognition Research Program',

Philosophy Compass, vol. 2 (2), pp. 338–346.

Simms, Eva-Marie (1996), 'Uncanny dolls: Images of Death in Rilke and Freud', *New Literary History*, vol. 27 (4), pp. 663-677.

Spaulding, Shannon (2010), 'Embodied Cognition and Mindreading', *Mind & Language*, vol. 25 (1), pp. 119–140.

— (2011), 'A Critique of Embodied Simulation', *Review of Philosophy and Psychology*, vol. 2, pp. 579-599.

Tenenbaum, Joshua B., *et alia* (2011), 'How to grow a mind: statistics, structure, and abstraction', *Science*, vol. 331, pp. 1279-1285.

Wimmer, Heinz and Perner, Josef (1983), 'Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception', *Cognition*, vol. 13 (1), pp. 103-128.

Zamir, Tzachi (2010), 'Puppets', *Critical Inquiry*, vol. 36 (3), pp. 386-409.