The Uncanny Valley Everywhere? On Privacy Perception and Expectation Management

Bibi van den Berg

Tilburg University,
Tilburg Institute for Law, Technology and Society (TILT),
P.O. Box 90153, 5000 LE Tilburg, The Netherlands

Abstract. In 1970 Mori introduced the notion of the ‘uncanny valley’ in robotics, expressing the eeriness humans may suddenly feel when confronted with robots with a very human-like appearance. I will use the model of the uncanny valley to speak about privacy relating to social network sites and emerging technologies. Using examples, I will argue that the uncanny valley effect is already manifesting itself in some social network sites. After that, I will project the uncanny valley into the near future in relation to emerging technologies, and argue that awareness of the uncanny valley effect is of great importance to technology designers, since it may be a factor in humans’ acceptance of and willingness to use these technologies. Will the uncanny valley be everywhere in the technological world of tomorrow?

Keywords: uncanny valley, privacy, social network sites, emerging technologies, robotics.

1 Introduction

In the spring of 2010 Google introduced a new social service called Buzz. Buzz is an add-on to Google’s GMail, which enables users to view information feeds about individuals they choose to follow, quite similar to Twitter. Days after its launch, there was a worldwide outcry regarding privacy violations in Buzz. What was all the buzz about? In an attempt to increase user-friendliness and ease of use, Google’s technology developers made a fundamental error. When users accessed their GMail in the days after Buzz’s launch, they were asked to complete a wizard to introduce them to this new service. In this wizard, Google’s technologists presented them with a list of individuals they were very likely to know. Buzz would make their GMail profile information would accessible to the individuals on this list, unless they opted out. Alternatively, users could sign up as followers of the individuals on the list, so that they would be kept up to date of changes in their profile information as well.

It turned out that users were collectively outraged by this automatically generated list of contacts. They claimed that encountering a list of individuals they

knew in an online setting that is generally considered to be private, such as one's e-mail facilities, made them concerned this was suddenly a *public* environment, and that, hence, their privacy had been violated. Moreover, since the lists of individuals they were presented with were quite accurate, many individuals felt highly uncomfortable and even scared by the realization of how much Google actually knew about them and the networks they operate in. What was worse, the lists were *quite* accurate, but not completely accurate. This made users feel eerie as well: how was it that Google knew about some of their social relations but not others, and how had the collection they were presented with – a mixture of well-known, intimate contacts and acquaintances that occupied only the fringes of their social circle – been compiled?

When discussing this example at an international conference in Texas this spring, danah boyd, one of the leading experts in research on teenagers' use of social media, made the following passing remark:

Google found the social equivalent of the uncanny valley. Graphics and AI folks know how eerie it is when an artificial human looks almost right but not quite. When Google gave people a list of the people they expected them to know, they were VERY close. This makes sense – they have lots of data about many users. But it wasn't quite perfect. [1]

Google Buzz's users experienced the setup of this new service as a privacy infringement and were left with an eerie feeling.

Some weeks after the Buzz buzz an 'eeriness incident' arose at my own institute. One bad day my boss, professor Ronald Leenes, received an e-mail from the social network site Facebook, supposedly sent by one of his contacts – whom for the sake of privacy we will call X – who is a member of Facebook. In this email X invited Ronald to become a member of Facebook as well. The e-mail did not just contain a message from X, but also a list of 'other people you may know on Facebook'. Note that Ronald himself is not a member of Facebook and never has been, and that the list of people presented to him, therefore, could not be based on Ronald's behaviors, address book, or on information he may have disclosed himself. It was based entirely on data about him that had been distributed by others – accidentally I presume. Using information about Ronald and his engagements with others, Facebook had built up a picture of his social network. What's more, eerily enough, the picture they presented was quite accurate. Ronald did indeed know almost all of these individuals. But as with the lists presented to users of Buzz, the picture was *quite* accurate, yet not entirely so. It was incomplete – some of Ronald's closest friends and colleagues were not on it, despite being active Facebook users –, and it was a bit of a haphazard collection: distant acquaintances were mixed with closer contacts. That made it even more eerie. Where did this collection come from and how had it been composed? How much does Facebook really know, not only about its users, but also about individuals *outside* the network?

These two incidents led me to start thinking about the term danah boyd used in the quote above: the 'uncanny valley', a concept well-known in robotics. I realized boyd had a point when she used this term in the context of social

media. Even more so, I realized that the uncanny valley may become a more frequently encountered phenomenon in the near future, when technologies will become increasingly autonomic and proactive. In this article I will explain why.

2 The Uncanny Valley – Theoretical Background

In a 1970 article the Japanese roboticist Masahiro Mori introduced the idea of an ‘uncanny valley’ in relation to the design of life-like and likable robots [2]. What did this valley consist of? Misselhorn summarizes the central thrust of the uncanny valley as follows:

... the more human-like a robot or another object is made, the more positive and empathetic emotional responses from human beings it will elicit. However, when a certain degree of likeness is reached, this function is interrupted brusquely, and responses, all of a sudden, become very repulsive. The function only begins to rise again when the object in question becomes almost indistinguishable from real humans. By then, the responses of the subjects approach empathy to real human beings. The emerging gap in the graph [...] is called the ‘uncanny valley’. The term ‘uncanny’ is used to express that the relevant objects do not just fail to elicit empathy, they even produce a sensation of eeriness. [3]

This trajectory is expressed in the image below. What does this picture tell us, exactly? On the far left of the figure, Mori says, we find ‘industrial robots’, such as robot arms. In designing these kinds of robots, which over the years have collectively come to be called ‘mechanical robots’ or ‘mechanoids’ [4], the focus is on functionality, rather than on appearance [2]. Robot arms don’t need to look like anything that provokes life-likeness or empathy in human beings – they need to complete specifically defined tasks, and that is all. Since functionality is the main design goal for these kinds of robots, mechanoids tend to be “*relatively machine-like in appearance*” [4].

However, there are also robot types that have more recognizable animal or even human forms. This group includes ‘toy robots’, which can be found halfway the rising line in Mori’s diagram, and ‘humanoid robots’, almost at the first peak. They are much more familiar to us and seem more life-like. Mori writes:

... if the designer of a toy robot puts importance on a robot’s appearance rather than its function, the robot will have a somewhat human-like appearance with a face, two arms, two legs, and a torso. This design lets children enjoy a sense of familiarity with the humanoid toy. [2]

This cluster of robots, which we will collectively call ‘humanoids’ is “*not realistically human-like in appearance*” and can easily be perceived as being robots, yet “*possess some human-like features, which are usually stylized, simplified or cartoon-like versions of human equivalents*” [4]. Sony’s robot dog AIBO and Honda’s humanoid Asimo are examples in case. With this type of robots we reach the first peak in the figure: they come across as being quite human-like and have a considerable degree of familiarity.

Fig. 1. The uncanny valley as introduced by Masahiro Mori in 1970

However, this is where an interesting turning-point emerges. If robots are *too* human-like, yet display behaviors that are less than perfectly human, Mori predicts there will be a steep decline in the level of familiarity, so much so that a deep sense of eeriness, or uncanniness, arises. If their appearance has “*high fidelity*”, Walters *et al.* write, “*even slight inconsistencies in behavior can have a powerful unsettling effect*” [4]. To explain how this works, Mori discusses the example of a prosthetic hand. He writes:

...recently prosthetic hands have improved greatly, and we cannot distinguish them from real hands at a glance. Some prosthetic hands attempt to simulate veins, muscles, tendons, finger nails, and finger prints, and their color resembles human pigmentation. [...] But this kind of prosthetic hand is too real and when we notice it is prosthetic, we have a sense of strangeness. So if we shake the hand, we are surprised by the lack of soft tissue and cold temperature. In this case, there is no longer a sense of familiarity. It is uncanny. In mathematical terms, strangeness can be represented by negative familiarity, so the prosthetic hand is at the bottom of the valley. So in this case, the appearance is quite human-like, but the familiarity is negative. This is the uncanny valley.[2]

Bryant remarks:

The uncanny valley itself is where dwell monsters, in the classic sense of the word. Frankenstein’s creation, the undead, the ingeniously twisted demons of *anime* and their inspirations from legend and myth, and indeed all the walking terrors and horrors of man’s imagining belong here. [5]

Why does the uncanny valley occur exactly? In general terms, one could argue that the sense of eeriness arises when a mismatch occurs between a robot's *appearance* and its actual *capabilities*. If a robot's looks are quite sophisticated, it seems logical that individuals interacting with it will assume that its behaviors will be quite sophisticated as well – just like we assume a level of behavioral sophistication from our fellow human beings whenever we engage in interactions with them. If there is a (significant) discrepancy between the high level of human-likeness of a robot's exterior and a low level of behavioral refinement, this provokes a response of disgust and repulsion: the uncanny valley.

At the right end of the diagram the valley is overcome. In the last stage of the figure the line goes up again. This is where we encounter interactions between real human beings: very human-like and very familiar. But some non-humans can also be placed on this rising slope, says Mori. These non-humans are not necessarily even *better* copies of human beings than the humanoids at the peak of the slope – quite the reverse is true. The solution of avoiding eeriness does not involve aspiring to ever more human-likeness, according to Mori. What is relevant is a greater degree of *behavioral* familiarity, or a greater similarity in *movements*. Mori uses bunraku puppets, used in the traditional Japanese puppet theater, as an example in case. These puppets are certainly not exceptionally human-like in their appearances – it is obvious for the viewers that puppets are used on the stage, rather than real human beings. However, Mori argues that since the movements of these puppets, generated by the human puppeteers working them, are remarkably human-like “*their familiarity is very high*” [2].

Based on the uncanny valley, Mori's conclusion is that robotics designers should limit themselves to creating relatively human-like, yet quite familiar robots. They should not aspire to create robots that mimic humans too closely. In the words of Bryant:

... designers of robots or prosthetics should not strive overly hard to duplicate human appearance, lest some seemingly minor flaw drop the hapless android or cyborg into the uncanny valley.[5]

In short: to avoid the emergence of the uncanny valley, according to Mori designers should focus on the creation of ‘mechanoids’ and ‘humanoids’ only. However, despite Mori's warning, a community of researchers has been conducting research into the creation of a third category of robots: ‘androids’. These are robots that have an “*appearance (and behavior) which is as close to a real human appearance as technically possible. The eventual aim is to create a robot which is perceived as fully human by humans...*” [4]. Whether their efforts will turn out only produce permanent inhabitants of the uncanny valley, or whether they will manage to get their machines to climb the final slope, time, practical constraints and our changing philosophical understanding of human-likeness will reveal. I will return to this point at the end of the article. For now, let's look at the impact that Mori's ideas have had in robotics and outside and see whether they can be applied to two new empirical contexts: that of social media and that of emerging technologies.

3 The Uncanny Valley Spills Over

After Mori's original Japanese article was finally translated into English in 2005 a surge of interest has emerged for this phenomenon. A number of empirical studies have been conducted in recent years to investigate whether Mori's theoretical paper could be underpinned with real-world evidence [6,7,8,9,10,11]. Some of these studies claim that Mori's model is (at best) an oversimplification of a complex world [6] or even that it does not exist at all [8]. Especially those working in the field of android development feel defensive about Mori's valley, and quite understandably so. If the uncanny valley does in fact exist, it threatens the viability of their research projects.

More refined (and less defensive) investigations regarding the empirical and conceptual basis of the uncanny valley exist as well [3,4,9,11]. For instance, some researchers argue that the valley may well exist, but can eventually be overcome in one of two ways. First, as time progresses robots will be developed that are ever more life-like, and display ever more complex behaviors, thus coming to mimic the behaviors and appearances of humans to an ever greater degree [12]. Second, some researchers argue that the occurrence of the uncanny valley is temporary in the sense that, as time progresses, human beings will become more and more used to dealing and interacting with (behaviorally and aesthetically) imperfect robots, and will thus overcome their initial sense(s) of repulsion.

Yet other researchers have aimed to uncover *what exactly* constitutes the eeriness of the valley, and *what exactly* causes it. For instance, Karl MacDorman argues that the uncanny valley "*elicits an innate fear of death and culturally supported defenses for coping with death's inevitability*" [11] – an explanation that is not wholly uncontested itself [3]. Mori himself argued that the valley emerged when individuals were confronted with life-likeness – with objects (corpses, prosthetic hands) that they assumed were 'alive' at first glance, and at closer look turned out to be lifeless. Lack of motion was key in the emergence of eeriness, he claimed. Other researchers have developed their own hypotheses on the appearance of the uncanny valley. For example, Hanson [7] has argued that the uncanny valley does not arise so much because of confrontation with too little *life-likeness*, but rather because of a confrontation with too little "*physical attractiveness or beauty*" [3]. One of the more convincing explanations, to my mind, is the idea that the uncanny valley effect has to do with our human tendency to *anthropomorphize* non-human and even non-living things, to attribute "*a human form, human characteristics, or human behavior to nonhuman things such as robots, computers and animals*" [13]. The more 'life-like' cues they give off, the more easily we will be tempted to ascribe intentions and animism to them, and the more easily we will be compelled to like them. All of this happens effortlessly, almost automatically, and mostly outside our awareness [13,14,15,16,17,18,19].

In the process of finding empirical and theoretical evidence for the uncanny valley, investigations into the matter have migrated from the field of robotics and android science into other research areas as well. For one, it has gained popularity, and found empirical proof, in the field of computer graphics and video game character design [10]. Moreover, it has been used in art analysis (e.g.

sculptures) and (animated) movie design. This is not surprising. As Bryant aptly summarizes it,

... though originally intended to provide insight into human psychological reactions to robotic design, the concept expressed by [the phrase ‘the uncanny valley’] is equally applicable to interactions with nearly any nonhuman entity.[5]

In the rest of this article I will attempt to apply the ideas underlying the uncanny valley to yet another domain: that of privacy perception in relation to (1) social network sites, and (2) emerging technologies, or to be more specific, to the proactive, autonomic technologies of the technological world of tomorrow.

4 Privacy Perception and the Uncanny Valley?

Let us return to the two examples discussed at the beginning of this chapter: the e-mail from Facebook, containing ‘individuals you may know’ and sent to a non-member of the network, and the automatically generated list of individuals to follow in Buzz, which led to widespread protests in the spring of 2010. What happened in both of these cases? Why did they cause a sense of eeriness, or, in the words of danah boyd, the “*social equivalent of the uncanny valley*” [1]?

One of the central ideas on the emergence of the uncanny valley and robots is that of *consistency*, or rather, a lack thereof. There ought to be consistency between the complexity and sophistication of robots’ *appearance* on the one hand, and the complexity and sophistication of their *behaviors* (movements) on the other. The moment this consistency is breached, a disconnect arises between what an individual expects of a robot in terms of behavior, based on its appearance and what he or she perceives in reality. This we have seen above.

Much of this relates to what is popularly known as ‘*expectation management*’. A robot that looks simple and highly mechanical evokes lower expectations in human beings in terms of refined and complex behavior than one that looks highly advanced and/or resembles human beings to a great degree. Now, let us generalize this principle to web 2.0 domains such as social network sites. When interacting with software, whether in internet environments or on our own computers, we also use a wide array of assumptions regarding what kinds of behaviors we may expect from these systems. Usually, the system’s appearance is a give-away in this respect: the more complex it looks, the more sophisticated its workings generally tend to be. In our interactions with various types of software, including online services, over the last years, we have thus built up a set of expectations with regard to the behaviors and possibilities of these systems relating to their appearance.¹ On the same note, systems we use for a wide array

¹ Of course, many more factors are relevant here, most notably our past use of these systems (or systems that resemble them or have a similar purpose or functionality). I am not claiming that appearance is the beginning and end-all of behavioral expectations with regard to software and online services. I am only pointing out that the system’s appearance is a relevant factor in a similar way that the appearance of a robot is a relevant factor in the way individuals predict their behavior.

of different tasks may reasonably be expected to be more complex in terms of the services they can deliver than environments in which we only conduct simple or limited tasks.

This is all the more true for add-ons. Most add-ons are single-purpose, simple little programs that add one limited feature to a system. And this is where the issues surrounding Buzz come in. Buzz was presented as a simple add-on, a new service to be added to users' GMail account, with one key functionality: keeping users updated on the information streams of individuals they chose to follow and making their own information stream public for followers in return. Buzz used a wizard that seemed accessible enough. By presenting individuals with a pre-populated list of possible followees, suddenly this seemingly simple and straightforward service displayed a level of intricacy and a depth of knowledge that caught users off guard entirely. Out of the blue Buzz displayed behaviors that most of us would consider quite sophisticated, and – even more eerie – it committed these acts with a seeming ease that surpassed even the most advanced computer users' expectations.

The Facebook e-mail is another example of the same mistake. It gathered a random collection of individuals that the receiver 'might know', and, what's worse, confronted a *non-member* of the network with this analysis of his social circle. It delivered no context and no explanation, but merely provided a simple, straightforward message. It left the receiver feeling eerie, precisely because the simplicity of the message greatly contrasted with the intricacy, the 'wicked intelligence' some would argue, of the system behind it. It revealed immense system depth – how else could such a complex deduction have been made? – but since the simple message displayed did not reveal anything about that depth, it left the receiver feeling uncanny indeed.

One thing is important to note. In robotics, we have seen, the uncanny valley emerges when the high level of sophistication of a robot's appearance does not match the limited level of sophistication of its behaviors. In these two examples the reverse is true: the *low* level of sophistication of the message/wizard does not match the immensely *high* level of sophistication of the behaviors suddenly displayed. In both cases, though, the disconnect between appearance and behavior is there, and hence an uncanny valley emerges.

Software developers and system designers need to take into consideration that the emergence of the uncanny valley is lurking in the corner when this disconnect arises. User-friendliness requires a close connection between behavior and appearance, so that users' expectations are properly met and eeriness, either because of overly simplistic behavior in a complex-looking machine or because of overly intricate behavior in a simple-looking machine, can be avoided.

5 Autonomic Technologies: The Uncanny Valley Everywhere?

In recent years visions such as autonomic computing [20,21,22], ubiquitous computing [23,24,25], the Internet of Things [26] and Ambient Intelligence [27,28,29]

have started predicting, and preparing for, a world in which technological artifacts will surround us everywhere and will constantly provide us with personalized, context-dependent, targeted information and entertainment services. They will not only do so *reactively* (i.e. in response to requests for information by users), but even *proactively*: technological artifacts and systems will cooperate to provide us with information and services that are relevant to us in the specific contexts in which we find ourselves, and they will do so *by anticipating what users' needs might be*. Users do not need to make their wishes explicit – the technologies in these visions of the technological world of tomorrow are said to be able to see users's needs coming, maybe even before the user himself knows. As I wrote elsewhere:

This aspect of Ambient Intelligence is by far the most far-reaching. It means, among other things, that systems will be given a large responsibility in managing and maintaining a user's information sphere. The technology [...] will decide what information is relevant, useful and even meaningful for the user in his current situation; the responsibility of finding, filtering and processing this information is removed from the user and placed squarely on the shoulders of the technology. It is the technology that will decide what is significant, and interesting, not the user. [29]

In order to be able to proactively provide users with the right kinds and types of information and services, technologies in tomorrow's world will use 'profiles', in which their preferences and past behaviors are stored, and thus, over time they will learn to adjust their behaviors to match users' situated needs and wants as perfectly as possible. The only way in which technologies in the world of tomorrow could learn these things is when they would have long-term, intimate contact with the user. One of the crucial questions that arises in relation to these technological paradigms, therefore, is whether (or to what extent) people are going to accept the fact that they will be watched and monitored by technologies always and everywhere, particularly knowing that all the information gathered thus is stored into profiles and used to make predictions with respect to future behaviors. As researchers in various fields have pointed out, there are serious privacy issues with respect to these types of technologies, since user profiles can act as infinitely rich sources of personal information [30,31].

But profiling may cause more than privacy problems alone. It raises a number of issues that could lead to the emergence of an 'uncanny valley' in users' perceptions of their interactions with technologies. First of all, one of the key parameters of technological visions of the future, such as Ambient Intelligence, the Internet of Things and ubiquitous computing, is the idea that technologies will be *embedded* into the background of our everyday surroundings – into walls, windows and floors in our homes, our offices and shops, and even in public spaces. Hiding technologies from view for users improves the ease of use of these technologies and makes them less obtrusive, which is especially important as more and more technological devices are gathered in our homes and workplaces, according to the originators of these visions [24,28].

However, the combination of embedding technologies into the background of everyday spaces *and* profiling users' behaviors to proactively provide them with appropriate and correct information and other services, leads to a high risk of ending up in an uncanny valley. After all, when technologies are hidden from view, how will users *know* whether they are being traced by technologies in any given situation, and how will they know what information about them is being captured and stored by cameras, sensors and other technologies in each situation? If technologies will proactively deliver personalized information, on the one hand, yet gather their data about users invisibly and imperceptibly on the other, chances are that users will regularly feel uneasy with the (eerily adequate) information they will suddenly be presented with, since they have no way of knowing how the system has come to collect that information, or how it has based its information provision on the users' behaviors. The discrepancy between the systems' complexity and the users' limited perception of that complexity (which is, after all, hidden from view) may lead users to feel the same way that non-users in Facebook or members of Buzz felt when confronted with a mismatch between a system's behavior (or system depth) and the same system's appearance: very eerie indeed. The mixture of proactivity, profiling and hiding technology from view may very well lead users' to feel unpleasantly surprised by the capacities and breadth of knowledge that systems may unexpectedly display in any given place or at any given time. As Abowd and Mynatt correctly remark, it is vital that users be aware of the fact that technologies are tracking them and how much data they are storing about them, lest they keep a sense of *control* over these technologies and accept them as part of their everyday life [32]. However, as said, the key design parameters of visions such as Ambient Intelligence and ubiquitous computing that we've discussed here – embedding technology, generating massive user profiles, acting proactively – appear to contradict that possibility and hence the emergence of uncanny valleys everywhere is a serious risk.

Second, users may find the personalized, targeted information thrown at them by these systems uncanny, because technological systems will act as though they know users intimately, yet because of the way profiles are composed, these systems are bound to make the wrong 'suggestions' every now and then, or deliver the wrong pieces of information. Profiling builds on combining two strands of information to create an expectation of individual users' future preferences, wishes and behaviors. I will use an example to show how this works. When online shoppers visit the bookstore Amazon.com, their past purchases and their search history are stored in an individual history. The past choices, wants and needs of each individual visitor to the store are registered, so that the store can provide him or her with ever more personally tailored product offers as time progresses and more visits have been brought to the store. At the same time, the behaviors of all the store's visitores combined are collectively aggregated, so that new visitors can find products matching their wishes more easily and efficiently. This is how that is done. Whenever a customer buys a book or a product, he or she will be offered a set of suggestions for products in the same category that were bought by other customers. The central idea is that "*if a group of people bought*

book A and also bought book B, others who might buy A might also be interested in B” [33] This form of so-called ‘collaborative filtering’ is also known as ‘planned serendipity’ [33].

In profiling these two streams of information are combined: the totality of past behaviors and choices of a single individual are merged with the collective behaviors of a large group of people, with respect to one single choice or purchase, and the integrated ‘image’ that arises thus is used to provide users with ever more accurate information and services, or so the thinking goes.

In many cases, in fact, this thinking is quite correct. The reason why a business such as Amazon.com deploys the profiling mechanisms that it uses, is – obviously – because they work: based on their buying behaviors one can conclude that shoppers are often quite happy with the product suggestions they are provided with. Needless to say, however, there will always be exceptions to the rule. Some buyers will feel uncomfortable with product suggestions as such, for instance because they will wonder how much Amazon.com is actually registering about them, or because they will feel uncomfortable that *some* of their past searching or buying behaviors come back to ‘haunt’ them in the store – imagine that you’ve conducted a search on a topic that also returned pornographic titles (which, of course, you were not out to find!), and being spammed with titles from that category ever after. Moreover, since people are different in their likes and dislikes, profiling based on the collective behaviors of a large group could easily lead to *quite* accurate product suggestions, but not entirely right ones. Compare this to the examples of Google Buzz and Facebook discussed above. Here, too, users may feel eerie because the personalized suggestions made are just a little too haphazard, and just a little too weird or unfit to match their exact preferences, although coming quite close... The uncanny valley lurks here as well.

Now, when realizing how central profiling mechanisms such as those used by Amazon.com will be in tomorrow’s world of smart, adaptive, proactive, personalized technologies, it becomes clear how urgent this issue really is. These two points raise the question: in the technological world of tomorrow, will the uncanny valley be everywhere?

6 Designers, Beware!

We have seen that Mori provided a solution for avoiding the eeriness of the uncanny valley. He argued that designers in the field of robotics ought to prevent the valley from opening up by limiting themselves to the creation of ‘mechanoids’ and ‘humanoids’, since the mismatch between appearance and behaviors was unlikely to occur in these types of robots [2]. Users wouldn’t expect a relatively simple-looking machine to conduct incredibly complicated tasks. Similarly, the more complicated the behaviors of a machine, the more refined its appearance ought to be as well – again, to ensure a match between humans’ expectations and reality. To my mind, the thrust of Mori’s argument was correct. At the time he wrote his article, very few actual examples of the uncanny valley effect existed in robotics – the science of robotics was simply not advanced enough for it to arise regularly. In the past decades, however, robotics has developed rapidly.

And, although their creators are sure to contest it, the uncanny valley does in fact arise in many who watch some of the creations that have been developed, especially with respect to androids.²

Moreover, as this article has shown, the uncanny valley effect may not necessarily be limited to robotics only. Perhaps the lesson that Mori's work teaches us is a broader one, which is that mismatches between *all* technologies' performances and their appearances ought to be avoided, if we strive for smooth interactions between humans and machines. Technology designers have a role to play in ensuring that a mismatch between impressions and expectations on the one hand and actual behaviors on the other does not occur. This does not only apply to robotics, but to any system we design. In a world that becomes ever more technologically saturated this is a lesson to be learnt sooner rather than later, lest we end up living with uncanny valleys everywhere.

References

1. boyd, d.: Making sense of privacy and publicity. Paper presented at SXSW in Austin (TX), USA (2010)
2. Mori, M.: The uncanny valley (translated by Karl F. MacDorman and Takashi Minato). *Energy* 7(4), 33–35 (1970)
3. Misselhorn, C.: Empathy with inanimate objects and the uncanny valley. *Minds & Machines* 19, 345–359 (2009)
4. Walters, M.L., Syrdal, D.S., Dautenhahn, K., Te Boekhorst, R., Koay, K.L.: Avoiding the uncanny valley: Robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion. *Autonomous Robots* 24(2), 159–178 (2008)
5. Bryant, D.: The uncanny valley: Why are monster-movie zombies so horrifying and talking animals so fascinating? (2004), <http://us.vclart.net/vcl/Authors/Catspaw-DTP-Services/valley.pdf>
6. Bartneck, C., Kanda, T., Ishiguro, H., Hagita, N.: My robotic doppelgänger: A critical look at the uncanny valley. Paper presented at the 18th IEEE International Symposium on Robot and Human Interactive Communication in Toyama, Japan (2009a)
7. Hanson, D.: Exploring the aesthetic range for humanoid robots. Paper presented at ICCS/Cog-Sci in Vancouver (BC), Canada (2006)
8. Hanson, D., Olney, A., Pereira, I.A., Zielke, M.: Upending the uncanny valley. Paper presented at the 20th National Conference on Artificial intelligence (AAAI) in Pittsburg (PA), USA (2005)
9. Brenton, H., Gillies, M., Ballin, D., Chatting, D.: The uncanny valley: Does it exist? Paper presented at the Conference of Human Computer Interaction: Workshop on Human-Animated Character Interaction (2005)
10. Schneider, E.: Exploring the uncanny valley with Japanese video game characters. Paper presented at the DiGRA 2007 Conference (Digital Games Research Association) (2007)

² For a personal experience of the uncanny valley, please watch (just a selection of many) 'eerie' clips of androids on YouTube: <http://www.youtube.com/watch?v=L4z3cs40cug>, <http://www.youtube.com/watch?v=MY8-sJS0W1I>, and <http://www.youtube.com/watch?v=091ugdijEM>.

11. MacDorman, K.F.: Androids as an experimental apparatus: Why is there an uncanny valley and can we exploit it? Paper presented at the Cognitive Science Society (CogSci 2005): Workshop 'Toward Social Mechanisms of Android Science' (2005a)
12. MacDorman, K.F., Minato, T., Shimada, M., Itakura, S., Cowley, S., en Ishiguro, H.: Assessing human likeness by eye contact in an android testbed. Paper presented at the Cognitive Science Society, CogSci 2005 (2005b)
13. Bartneck, C., Kulic, D., Croft, E., Zoghbi, S.: Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics* 1(1), 71–81 (2009)
14. Nass, C.I., Moon, Y.: Machines and mindlessness: Social responses to computers. *Journal of Social Issues* 56(1), 81–103 (2000)
15. Nass, C.I., Moon, Y., Fogg, B.J., Reeves, B., Dryer, D.C.: Can computer personalities be human personalities? *International Journal of Human-Computer Studies* 43(2), 223–239 (1995)
16. Turkle, S.: *The second self: Computers and the human spirit*. Simon and Schuster, New York (1984)
17. Turkle, S.: *Evocative objects: Things we think with*. MIT Press, Cambridge (2007)
18. Reeves, B., Nass, C.I.: *The media equation: How people treat computers, television, and new media like real people and places*. CSLI Publications/Cambridge University Press, Stanford (CA); New York, NY (1996)
19. Friedman, B., Kahn Jr., P.H., Hagman, J.: Hardware companions? What online AIBO discussion forums reveal about the human-robotic relationship. Paper presented at the Computer-Human Interaction (CHI) Conference 2003 in Ft. Lauderdale, FA (2003)
20. Ganek, A.G., Corbi, T.A.: The dawning age of the autonomic computing era. *IBM Systems Journal* 42(1), 5–19 (2003)
21. Sterritt, R., Parashar, M., Tianfield, H., Unland, R.: A concise introduction to autonomic computing. *Advanced Engineering Informatics* 19, 181–187 (2003)
22. Hildebrandt, M.: Technology and the end of law. In: Claes, E., Devroe, W., Keirsbilck, B. (eds.) *Facing the Limits of the Law*. Springer, Heidelberg (2009)
23. Weiser, M.: The computer for the 21st century. *Scientific American* 265(3), 66–76 (1991)
24. Weiser, M., Brown, J.S.: *The coming age of calm technology*. Xerox PARC, Palo Alto (1996)
25. Araya, A.A.: Questioning ubiquitous computing. In: *Proceedings of the 1995 Computer Science Conference*. ACM Press, New York (1995)
26. ITU: *The Internet of Things*. In: *ITU Internet Reports – Executive Summary: International Telecommunication Union* (2005)
27. Aarts, E., Harwig, R., Schuurmans, M.: Ambient Intelligence. In: Denning, P.J. (ed.) *The Invisible Future: The Seamless Integration of Technology into Everyday Life*. McGraw-Hill, New York (2002)
28. Aarts, E., Marzano, S.: *The new everyday: Views on Ambient Intelligence*. 010 Publishers, Rotterdam, The Netherlands (2003)
29. Van den Berg, B.: *The situated self: Identity in a world of Ambient Intelligence*. Wolf Legal Publishers, Nijmegen (2010)
30. Hildebrandt, M.: Defining profiling: A new type of knowledge? In: Hildebrandt, M., Gutwirth, S. (eds.) *Profiling the European Citizen: Cross-disciplinary Perspectives*. Springer Science, Heidelberg (2008)

31. Punie, Y.: The future of Ambient Intelligence in Europe: The need for more everyday life. *Communications & Strategies* 5, 141–165 (2005)
32. Abowd, G.D., Mynatt, E.D.: Charting past, present, and future research in ubiquitous computing. *ACM Transactions on Computer-Human Interaction* 7(1), 29–58 (2000)
33. Weinberger, D.: *Everything is miscellaneous: The power of the new digital disorder*. Times Books, New York (2007)