The Uncanny Valley

Daniel Mewes[1], Dr. Alexis Heloir[2]

[1] student at Saarland University, danielmewes@onlinehome.de
[2] researcher at DFKI Saarbrücken, alexis.heloir@dfki.de

Abstract. The uncanny valley is a phenomenon first described in 1970 by Masahiro Mori. It characterizes the correlation between the degree of human likeness of e.g. robots and their perceived familiarity by stating that, for a high degree of human likeness, subtle imperfections become more important and lead to an "uncanny" perception. This paper explains the notion of the uncanny valley and presents two different approaches of how to apply the phenomenon's consequences to the field of human computer interaction with a focus on robot design.

Keywords: uncanny valley, human likeness, robot design, social science research

1 Introduction

One of today's design goals for both humanoid robots and embodied agents it to let them appear as human like as possible. This goal arises from the idea that a high level of human likeness ensures that e.g. a robot gets easily accepted by humans and allows for a natural and rich way of interaction. Human like robots allow for the expression of emotion and the usage of both subtle and well visible gestures, like pointing at objects or focusing with their eyes. They can therefore utilize the natural way of inter-human communication, which is not limited to speech but apparently also incorporates a good part of body language and facial expression.

In many areas this rich way of interaction is actually vital. For example somewhere in the future robots might play a role in caring for children, elders or sick people. Especially for childcare a human like appearance and a natural way of interaction might be extremely important for pedagogical reasons. Also a small child cannot be given a complicate user manual that describes how it has or has not to interact with some robot.

As of today, human like robots are used for product presentations on exhibition booths. Their human like appearance allows for an increased attraction of customers to the booth.

On the scale from not human like at all to very human like we distinguish between industry robots, humanoid robots (or humanoids) and androids (see Fig. 1). While industry robots are designed to be purely functional and human likeness or familiarity is not a goal in their design at all, humanoid robots are designed in order to resemble the basic shape of human beings. They usually feature legs, arms, a torso and head. Many of them also have two independent eyes or a similar structure and some equivalence of a mouth that may be movable for a more natural appearance. Androids are designed with the goal of ultimately becoming indistinguishable from real humans. They typically feature a human like skin, artificial hair and a face that is capable of facial expression and have the exact proportions of a human (typically chosen to meet some ideal of beauty). While humanoid robots are often perceived as being neuter, androids are clearly distinguishable as being either female or male.

Fig. 1. Different degrees of human likeness. From left to right: Industry robot, humanoid, android

1.1 The Uncanny Valley

The uncanny valley as described by Mori [1] is a phenomenon that arises when setting human likeness in relation to familiarity. A high familiarity rating of something includes that humans can quickly build up some kind of emotional relation to the thing and may perceive it as being trustworthy. It also induces the instant possibility of interacting with the thing rather than first having to become accustomed to it. In Mori's 1970 text, the term familiarity is also extended onto the negative scale. According to Mori, negative familiarity is equivalent to uncanniness or eeriness, although this equivalence is considered controversial especially from a psychological point of view (compare [2] p.304).

One's intuition may tell that with an increase of human likeness, there also comes increased familiarity. Unfortunately this only seems to be true for a limited range of human likeness. The “uncanny valley” phenomenon kicks in, whenever a high but not yet perfect degree of human likeness is reached. At some point, familiarity steeply breaks down when further increasing the human likeness. It then reaches a degree of negative familiarity which even lies below the familiarity rating of non human like things. Only when human likeness approaches perfection, familiarity rises again and ultimately reaches a level where it is on-par with the perceived familiarity of real humans (because of the thing being actually indistinguishable from those at that point). The resulting shape of the human-likeness/familiarity graph as depicted in figure 2 is what gives the “uncanny valley” its name. One should note that figure 2 does not show an exact quantization of both axes however. An attempt of quantization of the uncanny valley was made in a study described in [2] pp. 304-309, which we are going to present in more detail in section 2.1 (optical appearance).

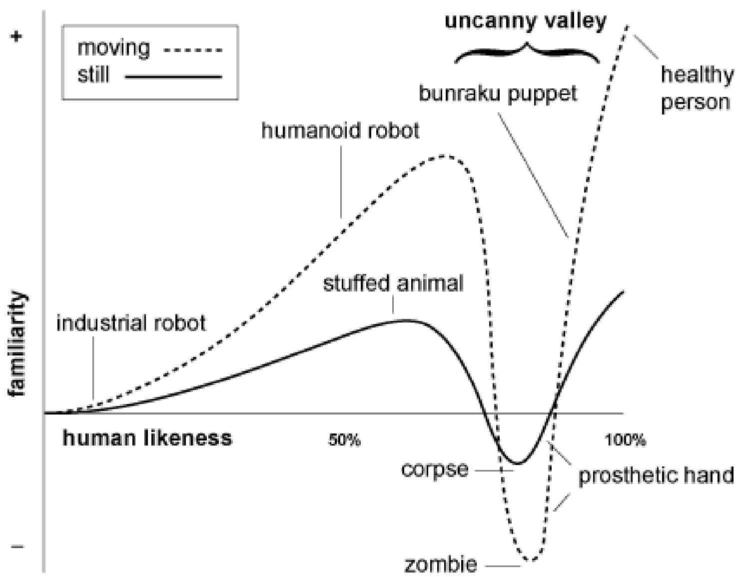

Fig. 2. The uncanny valley graph from Mori's article in Energy (simplified)

1.2 Refinement of the Uncanny Graph

The actual perception of a robot or embodied agent with respect to its familiarity does not depend on its human likeness alone, but also on the time of exposure and the degree of expression and interaction. Mori suggests that human like things which are moving, amplify the feeling of both familiarity as well as uncanniness. The refined familiarity/uncanniness graph for moving things in included as a dotted line in figure 2. While according to this, static non-moving objects can never reach a really high level of familiarity, movement also increases the risk of “falling” into the uncanny valley. Many androids look very natural and may even be considered as indistinguishable from humans when seen on a photograph. But if they start moving around, many are perceived eerie after a sufficient time of observation. On the other hand, according to a study published in [3], only 23% of the participants considered a static android as being human, while 70% did so for a moving but otherwise identical android in a two seconds test. The time of exposure is another important factor in familiarity considerations. This is obviously true because of the fact that it is much easier for an android or agent to maintain the illusion of perfect human likeness if observed for a short time only, since there just is less "opportunity" for it to look or move strangely.

The idea behind the refinement of the uncanny graph regarding movement may also be extrapolated to interactivity. An android for example may look and move extremely human like, making it very familiar. If we start interacting with it, we may however quickly notice that it reacts strangely, which may eventually lead to an uncanny impression. Mori describes a similar phenomenon in his 1970 paper [2]:

> "[R]ecently prosthetic hands have improved greatly, and we cannot distinguish them from real hands at a glance. [...] So maybe the prosthetic arm has achieved a degree of human verisimilitude on par with false teeth. 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 this example it is touching the hand that suddenly makes it become uncanny, although limiting the degree of interaction to visual observation allows the maintenance of a nearly perfect life-like perception.

2 Application to Human Computer Interaction

The uncanny valley phenomenon can play a role whenever deploying androids or human like embodied agents as means of human computer interaction. Because of the potential problems arising from it, it seems reasonable to consider the phenomenon at the design stage. We first cover the approach of an incremental design process, which, aside from showing up a way of dealing with uncanniness problems, also provides a foundation for research of social behavior and human perception. We also present an alternative approach to circumvent problems possibly arising from the uncanny valley phenomenon.

2.1 Incremental Design

In order to overcome the uncanny perception of a human like android, we may want to further improve its human likeness. To achieve this, we have to know how the human perception of androids works and what features are necessary in order to be perceived as human like. One could therefore guess about those features or try to obtain them by observing the appearance, behavior and interaction of human beings. The efficiency of this approach is at best arguable however. Therefore, MacDorman and Ishiguro (in [4]) propose the idea of an incremental design, which allows incrementally improving upon the design of an android and at the same time gaining new insights about human perception.

The basic idea of the incremental design approach is to start by building an actual android based on "best guesses" or existing knowledge about how an android should look, move and behave. Although probably not perfect, this android can be used to perform studies regarding different aspects of social science and human perception. These studies may lead to new insights about how the human perception of androids works and which parameters and aspects must be changed or redesigned in order to improve upon the perceived human likeness of the android.

Depending on the studies, it may even be possible to gain knowledge about how exactly some parameter has to be set for an optimal design. For example Nancy Etcoff reports in her 1999 paper [5], that changes in the scale of millimeters of the eye-to-eye distance of a face can be crucial for whether the face is perceived as attractive or ugly. MacDorman and Ishiguro propose that experiments to find out such parameters may be a good application for android based studies, since these parameters "could be manipulated with an android, but not a human actor" ([2] p. 315).

After all studies have been carried out, the next step in the incremental design approach is to utilize the insights from the studies in order to improve upon the design of the android. Either the existing android can then be altered accordingly or a new improved one can be built. Since the biggest problems of the previous design can be sorted out this way, the improved android may then be used for new studies, in order to get insights about more subtle problems. These new studies may not have been possible with the first design because of the subtler problems being "covered" by the bigger problems that existed. The procedure can then be repeated as long as required (Fig. 3).

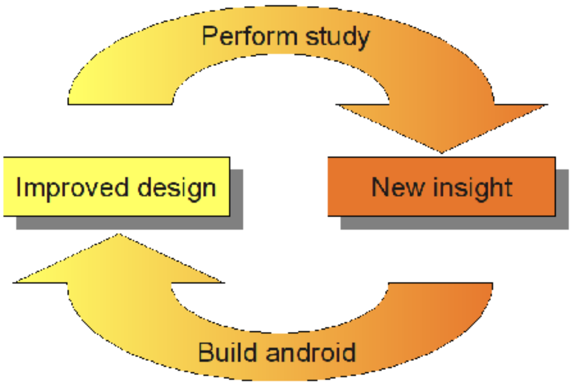

Fig. 3. Circle of incremental design

Areas of Incremental Design

There are different areas for which the incremental design approach can be deployed. We present three of them – optical appearance, movement and interaction – in this section. In general it is important to note that especially the areas of optical appearance and movement may intersect with each other in experiments, making it hard if not impossible to improve the overall design of an android by just focusing on one of those. We provide an example for such an experiment in the section about movement.

The field of interaction on the other hand can vary in its importance, depending on the intended use of the android. Also many experiments regarding appearance and/or movement are possible without any kind of interactivity, which sometimes allows for an isolated approach that does not include interaction.

Optical Appearance

Not only is a good appearance obviously an important aspect for human likeness, but MacDorman and Ishiguro also state that there is an advantage of deploying androids in appearance related studies in order to gain a better understanding of interpersonal responses between humans. "[...] it is nearly impossible for a human actor to hold a static pose. Out bodies are constantly moving, however slightly" ([2] p. 314) they are saying. Therefore androids allow for a much more precise and isolated examination of the specific area of appearance. As stated before regarding a possible eye-to-eye distance experiment, androids also allow for certain aspects of their appearance to be altered freely, which cannot be achieved using actual human beings. Even if one can utilize a variety of differently looking human models for an experiment, their appearance will certainly not be the only difference between those and additional seemingly unrelated aspects may have an influence on an experiment's outcome. With androids on the other hand, the appearance can be altered completely orthogonal to all other aspects (like movement for example).

An experiment described in [2] gives an idea of how an appearance related android based study can be designed. In the experiment, 45 Indonesian participants were shown a total of 31 different pictures in random order. Each of the pictures depicted either the face of an humanoid robot, an android or a human, or some mixture of humanoid and android faces or android and human faces in different ratios. The participants were asked after each picture to rate their impression regarding the factors human likeness, familiarity and eeriness. The outcome of the experiment can be seen in figure 4 (only a selection of the pictures is shown). This specific experiments actually provides a quantization of the uncanny valley phenomenon. Similar experiments could be deployed in order to compare different android designs or aspects of those regarding their effect on the mentioned (and possibly additional) factors.

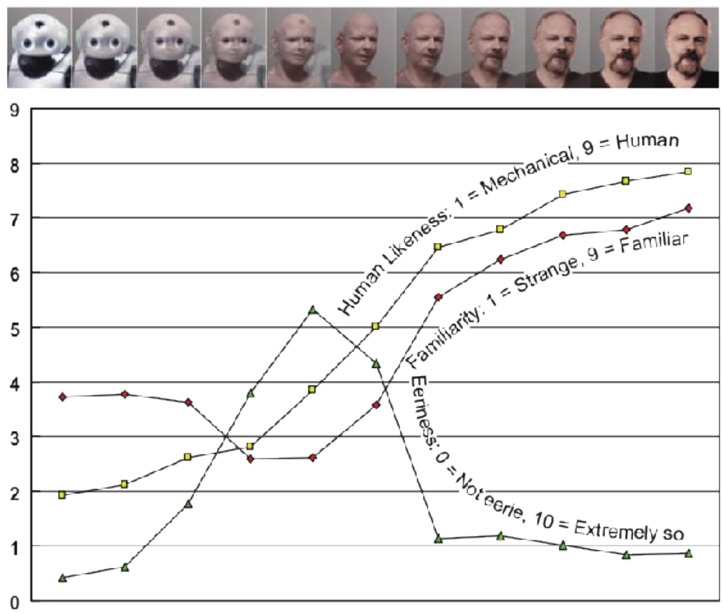

Fig. 4. The face-morphing experiment reveals how different faces are perceived by participants regarding their familiarity, human likeness and eeriness

Movement

It seems that the movement of an android is absolutely vital for perceiving it as being human-like. As already mentioned earlier, in 2005 Hiroshi Ishiguro published the result of an experiment (compare [3]) in which the participants were shown an android for two seconds and were asked to tell whether what they saw was a human being or an android. The experiment was performed in two different ways. Some participants were presented a static non-moving android and some a moving one. The static android was considered human by 23% of the participants, while the moving one was so by 70%.

In an experiment by Minato and others (compare [6]), eye-tracking was used while participants were set to converse with either a human child or an android child. The experiment revealed that when conversing with a human, the participants mostly fixated on the right eye of the child. Interestingly, the participants showed the same fixation pattern when conversing with the child android. The experiment then was repeated with a mechanically looking but equally gesturing android. This time the eye-fixation patterns were different, the participants now “were treating the humanoid robot more like an object than a person, since they spent more time looking at parts of the robot's body other than the eyes” ([2] p. 316). What this experiment shows is that although the movement of an android is vital for human likeness, it's appearance is important too and it does not seem to suffice to focus on either the one or the other only, when designing or improving an android.

Interaction

The area of interaction is a huge field. Starting with subtle but nonetheless important capabilities like holding eye contact to a person or reacting on other people's gestures, it goes all the way up to holding meaningful conversations and other similarly complex tasks. It is clear that a perfect human like interactive android therefore is extremely hard to built, at least with today's generally available technology. There are however aspects were some limited capabilities for interaction can be incorporated into androids to vastly improve upon their human likeness.

Repliee Q2 by Kokoro Company for example has special sensors to locate people in its environment in order to turn its head into their direction and hold eye-contact with another person. Q2 was actually shown performing an interview, where this capability plays an important role. Another example for interactivity is the Geminoid android – also by Kokoro Company – which turns its head away and changes its facial expression when touched near the neck.

Using androids for interactive studies allows for interesting insights into humans' interpersonal responses, which regarding to MacDorman and Ishiguro “would be far more difficult, if not impossible, to achieve [...] with human actors”. Namely, an eye-tracking study performed by MacDorman and others ([7]) showed that “Japanese participants tended to break eye contact when asked questions that required thinking” ([2] p. 316) and typically look down while thinking. The interesting result then is that this focus pattern was only observable when the participants were conversing either to human beings or with androids which they thought were under human control. When believing that the android was under its own control, the downward pattern disappeared. The experiment therefore reveals that the eye fixation behavior of (at least Japanese) humans in interpersonal interaction depends on what a person thinks about the mind of their dialog partner.

An Example of Incremental Design

The “Actroid” line of robots (see Fig. 5) by Kokoro Company Ltd. is an excellent example of incremental design at work. In January 2004 (compare [8]), Kokoro Company presented the “Actroid Repliee Q1”, an android especially built for research purposes. Repliee Q1 was used in a series of studies in order to gain better knowledge about human perception and social mechanisms. The androids design was then improved and in December 2004 came the “Actroid Repliee Q1-Expo” with more capabilities and a higher degree of human likeness. After further research including studies with Repliee Q1-Expo, the next successor “Actroid Repliee Q2” was revealed in July 2005. Repliee Q2 again serves the purpose of helping in research. The previously mentioned “Geminoid” android also is an offspring of the Acroid line built for research purposes.

In addition to the Repliee line of androids, Kokoro Company utilizes the knowledge and insights gained through their research for commercial androids. At the “Expo 2005 Aichi Japan”, Kokoro demonstrated “Actroid-expo”, which aims at being deployed as a robot receptionist. Another line is constituted by the “Actroid DER” androids, where “DER” is an acronym for “Dramatic Entertainment Robot” (according to Wikipedia [9]). Aimed at amusement, presentations and shows, the Actroid DER androids are available for rental by customers. After the release of the first “Actroid DER” in June 2005 came the improved “Actroid DER2” in October 2006 and the further improved “Actroid DER3” in October 2008.

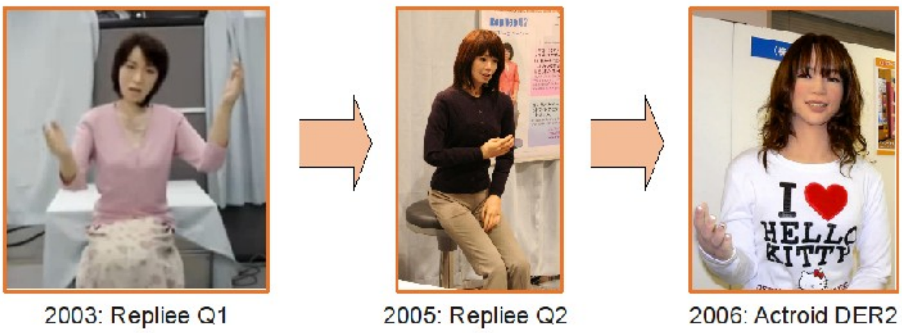

Fig. 5. Androids from the “Actroid” line

2.2 Moderate human likeness approach

While the incremental approach is useful for gaining insights for social science and to improve upon the human likeness ob androids, it may not always be the right approach for maximizing the familiarity of robots. Because of the uncanny valley phenomenon, very human like androids have a high probability of becoming eerie and therefore unfamiliar.

Mori in his 1970 Energy article [1] was aware of this consideration and supposed an "escape by design" in order to circumvent it. Although technical possibilities have somewhat improved since 1970, Mori actually bases his considerations on prosthetic hands, which had already been well developed back then. Mori says that "recently prosthetic hands have improved greatly, and we cannot distinguish them from real hands at a glance. [...] If we shake the hand, we are surprised by the lack of soft tissue and cold temperature. [...] It is uncanny." ([1]) and these or similar problems still apply to today's advanced androids. On the other hand, Mori refers to glasses and wooden hands without joins which do not exactly resemble real eyeballs respectively human hands, but are still considered "charming" or "beautiful" ([1]) without a sense of uncanniness. He recommends "designers [to] take the first peak as the goal in building robots rather than the second", referring to the peaks in familiarity that can be seen in the uncanny valley graph (compare figure 2).

In 2009, Japanese researchers under Shuji Kajita presented the "HRP-4C" "fashion bot" (see figure 6). According to an interview (see [10]), the designers "deliberately leaned toward an anime style" in order to circumvent the uncanny valley. Although they may not have been very successful with this and actually revised the design shortly after the first presentation in order to further decrease the human likeness of the robot (see [11]), this example shows that the "escape by design" approach actually plays a role in real world robot design.

Fig. 6. HRP-4C "fashion bot" has been designed with limited human likeness in mind

3. Conclusion and Evaluation

Building human like robots has both practical as well as scientific advantages. By deploying androids in social science experiments, new insights can be gained about how to achieve a better degree of human likeness. Such studies may also provide insights into human perception and social behavior and should therefore be considered whenever humans are unable to provide the required degree of exactness and/or orthogonality between the relevant factors for some experiment.

Depending on the exact design goals (id est maximum familiarity versus maximum human likeness), one may either consider the approach of an incremental design or the intentional limitation of the robot's human likeness.

The different approaches to familiar robots can be observed in Hollywood features, too. An example for maximal human likeness is Steven Spielberg's “AI – artificial intelligence”. In this movie, the protagonist is an extremely human like child android which was built with emotions and consciousness. While identifying with the protagonist is easy, there are a few scenes in the movie where the artificial nature of it becomes visible and I (Daniel Mewes) personally find these scenes to be very uncomfortable. In the movie “I Robot” on the other hand, one of the protagonists is a human like but still mechanical looking robot. Although identifying oneself with it may be a bit harder at first, developing a good amount of sympathy appears to be possible anyway and there absolutely are no uncanny scenes with this robot. The moderate human likeness approach may therefore be favorable for many applications.

References

1. Masahiro Mori: The Uncanny Valley. Energy 7(4) pp. 33-35 (1970), translation by Karl F. MacDorman and Takashi Minato
2. Karl F. MacDorman, Hiroshi Ishiguro: The uncanny advantage of using androids in cognitive and social science research. Interaction Studies 7:3 (2006)
3. Hiroshi Ishiguro: Android science: Toward a new cross-disciplinary framework. Cog-sci-2005 workshop: Toward social mechanisms of android science, Stresa, Italy (2005)
4. Karl F. MacDorman, Hiroshi Ishiguro: Toward social mechanisms of android science: A CogSci 2005 workshop. Interaction Studies 7 (2006)
5. Nancy Etcoff: Survival of the prettiest: The science of beauty. New York: Doubleday (1999)
6. T. Minato, M. Shimada, H. Ishiguro, S. Itakura: Development of an android robot for studying human-robot interaction. In: R. Orchard, C. Yang and M. Ali: Innovations in applied artificial intelligence, Springer-Verlag Berlin (2004)
7. K. MacDorman, T. Minato, M. Shimada, S. Itakura, S. Cowley, H. Ishiguro: Assessing human likeness by eye contact in an android testbed. Proceedings of the XXVII annual meeting of the cognitive science society, Stresa, Italy (2005)
8. Kokoro Company Ltd.: Actroid History http://www.kokoro-dreams.co.jp/english/robot/act/actroid_rekisi.html
9. Various: Wikipedia article on “Actroid” as of Jul. 23rd 2009 http://en.wikipedia.org/w/index.php?title=Actroid&oldid=303750074
10. ABC News: Fashion fembot to strut Japanese catwalk (2009) http://www.abc.net.au/news/stories/2009/03/16/2517770.htm
11. “Taro”: Revisiting Japan's 'Uncanny Valley' http://news.3yen.com/2009-03-24/revisiting-japans-uncanny-valley/