

Spooked?

What is it about almost-humans that gives us the creeps, wonders Laura Spinney

ON THE face of it, *The Polar Express* was a sure-fire winner: starring Tom Hanks, it told the charming story of a boy's magical train journey to the North Pole. But when the movie came out in 2004, there was a problem: the ultra-realistic animation gave some viewers the creeps. Five years later, when James Cameron chose the same technology for *Avatar*, his graphics people reportedly thought the decision might bankrupt the production company. But Cameron's blue humanoids went down a storm. For a while, *Avatar* was the highest-grossing film of all time.

You might have heard of the uncanny valley: the notion that the more human-like a non-human character becomes, the more we like it – until suddenly, we don't. At some point where it is almost, but not quite, human we become unsettled, even revolted. The uncanny valley has been used to explain our adverse reactions to all sorts of almost-humans from zombies, androids and corpses to the creepy clowns recently terrorising North America. The characters in *The Polar Express* strayed into the valley. Cameron's blue Na'vi did not. Why?

The Japanese roboticist Masahiro Mori first described the uncanny valley in 1970. It has since become a highly influential idea that has shaped not only films and video games, but also robots, dolls and prosthetic limbs. But it has only recently been tested scientifically. The result? Researchers are divided on almost every aspect of it, from why we experience it to

whether it actually exists. The uncanny valley, it seems, is weirder and more controversial than Mori could have predicted.

Mori's original paper was a warning to roboticists not to stray too close to human likeness, or risk repulsing the very people they wanted their robots to serve (see "The unbearable likeness of beings", page 31). But his thought experiment was based on intuition rather than empirical research, and evidence has been equivocal. This has led some to suggest that Mori was wrong. "The uncanny valley doesn't exist," says David

"Robots that look too human risk repulsing the people they are meant to serve"

Hanson at Hanson Robotics in Hong Kong. He agrees that human-like figures can elicit an uncanny or eerie sensation, but sees no evidence for a "valley" that opens up at a certain level of human resemblance.

Others believe that the valley does exist, but not as generally interpreted. "It's not meant to be literally true," says Karl MacDorman, who studies human-computer interaction at Indiana University in Indianapolis. "It's about the risk of provoking certain reactions." Understood correctly, he says, the uncanny valley can aid designers and even provide a tool for understanding dealings between humans, by comparing them with

human-android interactions.

Part of the problem, MacDorman says, is that Mori's meaning was lost in translation. Writing in Japanese, Mori coined the word *shinwakan* to describe our emotional response to a human-like creature. Over the years, *shinwakan* has been translated variously as affinity, familiarity, rapport, comfort level and likeability. These all have slightly different meanings, and crucially, some can have negative or positive values – likeability, for example – whereas others, such as familiarity, can't dip below zero. This might explain, among other results, why neuropsychologist Marcus Cheetham of the University of Zurich, Switzerland, and his team failed to find the uncanny valley in 2014 when they measured people's responses in terms of familiarity.

And there's a second reason why the phenomenon might have proved elusive: lab tests have often used computer-generated morphs that gradually turn a robot into a human. Volunteers are aware that these "hu-bots" cannot exist in the real world, and this may affect how they respond to them. Earlier this year, Maya Mathur, a cognitive scientist at Stanford University in California, and her colleague David Reichling of the University of California, San Francisco, tried to solve this problem by showing people images of 80 actual robots. Their findings support the uncanny valley and, says Mathur, give a crude indication of its location and extent (see "Locating the uncanny valley", page 30). ➤



LUISA WHITTON

Among those who believe that the valley exists, there are two broad schools of thought as to its cause. Both tap into our brain's evolved ability to detect a threat, and hence to survive in a dangerous environment such as the African savannah. The first puts it down to a phenomenon called category uncertainty. When we are unsure how to categorise a stimulus – is it lion or antelope, human or machine – that ambiguity gives rise to fear, which in turn triggers the fight-or-flight response. The second school invokes something called perceptual mismatch, which occurs when an entity has seemingly incongruous features. The mouth and nose may be human-like, for example, but the eyes appear dead.

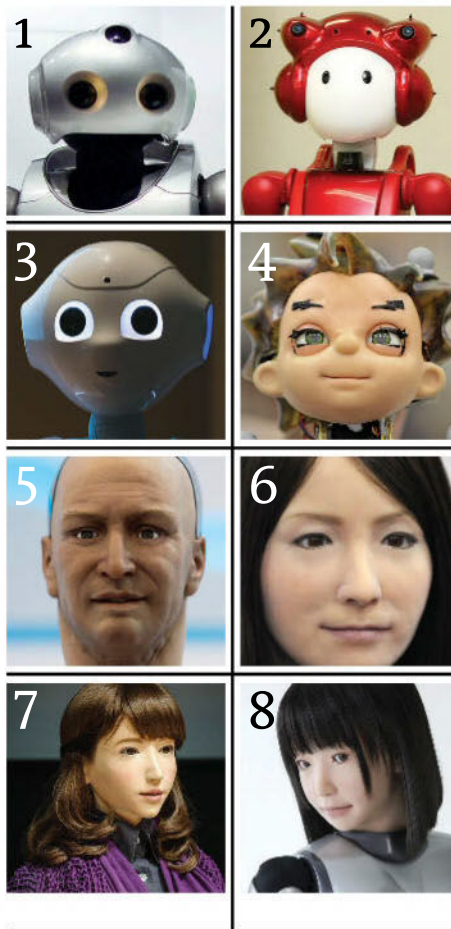
Whereas category uncertainty is conscious, perceptual mismatch is unconscious or preconscious. Your brain detects a conflict and infers a threat, but cannot consciously identify its source – perhaps something corpse-like about the creature taps into a hardwired desire

“Uncanniness taps into our brain’s evolved ability to detect a threat”

to protect ourselves from a potential source of infection, for instance. “Instead of producing fear, it produces anxiety – a kind of free-floating anxiety that can’t attach itself to any particular cause,” says MacDorman, who came up with the perceptual mismatch model. It’s an adaptive emotional response that enables us to maintain vigilance in the face of uncertainty.

The two explanations have very different implications. If category uncertainty is correct, to avoid eeriness roboticists and others should try not to make their creations too human-like, as Mori advised. In other words, they need to steer clear of the human-robot boundary. But if perceptual mismatch is the cause of uncanniness, avoiding it is more about maintaining consistency within the entity. Say you can give an avatar a highly realistic skin texture, for example, but truly lifelike eye movements are beyond your technical ability. To prevent your creation stumbling into the valley, you would deliberately use more artificial-looking skin so that the overall effect is a consistent level of realism.

Which idea is correct? Last year, a team of cognitive and computer scientists at Aalto University in Espoo, Finland, led by Jari Kätsyri, reviewed previous studies to see how well each hypothesis was supported. Starting with a sample of 125 peer-reviewed papers,



LOCATING THE UNCANNY VALLEY

Humanoid images generally become more likeable as they become more humanlike (from top left). But according to research by Maya Mathur at Stanford University, somewhere around image 4 a sense of unease sets in for most of us - before disappearing again around image 6

they removed those that didn't meet their standards – not having enough volunteers, for example, or including morphed images of hu-bots that could not possibly exist in real life. They ended up with just 17 studies.

Overall, these did not support Mori's original concept, tending instead to show a simple positive correlation between likeability and human likeness. However, they did show that the uncanny valley could arise under certain conditions – notably, those of

perceptual mismatch – and that incongruous eyes were particularly responsible for conjuring up eeriness. The Finnish group drew one more conclusion: corpses and zombies, which occupy the deepest part of the valley in Mori's graph, should be excluded from the debate. “They are a distraction, because they elicit uneasiness for reasons besides being human-like,” says Kätsyri.

Soon after that review appeared, MacDorman and Debaleena Chattopadhyay, now at the University of Illinois, Chicago, published a study that pitted the category uncertainty and perceptual mismatch models against each other. They took a human face and gradually transformed it into a computer-animated version of itself. When they showed this to 500 volunteers, they found that people had no trouble categorising the fully computer-animated face as an animation, but they also found it the most eerie. In other words, although category uncertainty was not at play, there was still uncanniness. And the feeling of uncanniness could be enhanced if they made certain features within the morphing face – especially the eyes and mouth – appear more animated than the rest. MacDorman and Chattopadhyay concluded that perceptual mismatch was the most likely explanation for the uncanny valley.

The perceptual mismatch theory is much better investigated than category uncertainty, however, and some researchers believe the latter remains a contender. Psychologist Tyler Burleigh at the University of Guelph, Canada, thinks MacDorman and Chattopadhyay found no category effect because they did not test actual category boundaries – whether real or animated, the faces were still human. In 2013, Burleigh's team showed people two series of images, one depicting the transformation from a human to a goat, the other from a human to an android-like creature with grey, inhuman skin. In other words, they toyed with the basic categories of human, animal and machine. “That’s where we find the effect – specifically where we manipulate the category membership,” says Burleigh.

So, it looks as if the jury is still out. And the questions don't end there. Back in 1970, Mori wrote that the way a creature moves is critical to invoking eeriness, but very few researchers have explored the effect of movement. One who has is cognitive scientist Ayse Saygin at the University of California, San Diego. She has used functional MRI to show that people's brains register an error when they see a human-like creature moving robotically, but not when they see either a human moving



While some may want to harness such reactions, in other domains the uncanny valley is to be avoided at all costs: cosmetic surgery, for one. At the University of Louisville in Kentucky, plastic surgeon Joshua Choo is designing a study to explore when surgically

“In cosmetic surgery, the uncanny valley must be avoided at all costs”

altered faces trigger perceptual mismatch. “A patient will come to you and say, I want this person’s nose,” he says. “And you can have a successful result, if you look at the nose and crop out everything else, but when you take it in the context of the whole face, it looks odd.” Choo hopes his findings will help those in his profession avoid creating that “operated-on” look.

Others who work on the margins of the uncanny valley and would rather not stray inside have a different problem. Kätšyri, now at Maastricht University in the Netherlands, thinks that if the perceptual mismatch theory is correct, the uncanny valley will be increasingly difficult to avoid as robots and avatars become more realistic, because some elements of their appearance just won’t be as polished as others. That doesn’t mean designers should shy away from what he calls the “grand challenge” of building human-like beings. But for many practical applications, there may be wisdom in Mori’s advice to go for only moderately human-like androids.

Nevertheless, some androids are already so realistic that they seem to have escaped the valley and regained the sunny uplands of likeability again – at least in static images. These include several built by Hanson’s company. Being sceptical of the uncanny valley phenomenon, Hanson finds it easy to ignore it when designing robots, but he thinks that the very idea of it might be limiting creativity in other designers by discouraging them from exploring how to get around or even harness uncanniness.

Perhaps fittingly, MacDorman takes the contrary view: that only by understanding the uncanny valley can you “design your way out of it”. Either way, we still have more to learn if we want to harness the power of uncanniness – whether it be to create blockbusting movies with avatars or to keep spooky almost-humans at bay. ■

naturally, or a robot moving mechanically. Kätšyri even suggests that there may be more than one uncanny valley, occupying different dimensions of perceptual experience. It’s possible, for example, that some entities do not evoke uncanniness until they move.

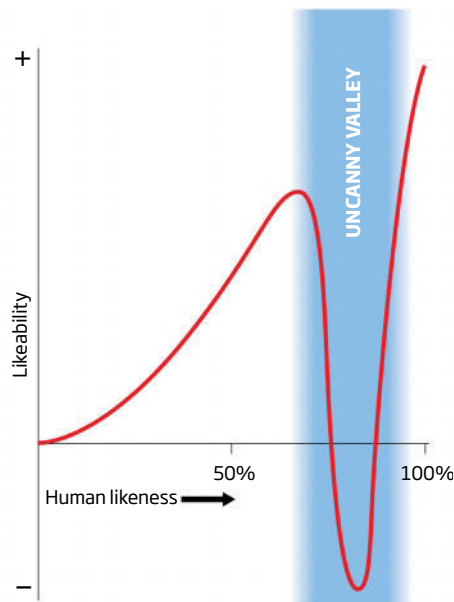
If the cause of uneasiness is still up for grabs, our reactions to it are even less well understood. Oddly, we don’t necessarily run from the uncanny valley. Sometimes we are drawn to it, showing a morbid fascination for the eerie and the almost-human. The entire horror film industry depends on this reaction, and video games are now capitalising on it too. Burleigh cites the popular *Silent Hill* series as an example: “There’s a really good case of using human-like monstrosities to enhance the horror experience,” he says. However, to exploit this effect, it would be useful to understand the circumstances in which uncanniness elicits morbid fascination or another paradoxical response, amusement.

The first ever study of uncanniness and humour surfaced only last year, when Meeri Mäkäräinen of Aalto University presented a paper entitled “The Funcanny Valley” at a computer science conference in Finland. Working with Kätšyri and others, she discovered that exaggerating a person’s smile beyond realistic dimensions did not necessarily elicit a negative reaction. “When our participants saw strange faces, they smiled instead of frowning,” she says. Of course, caricaturists and the makers of latex puppets of politicians have been aware of this for some time. Clowns may be unwittingly cashing in on it too – although coulrophobia,

or fear of clowns, could mark the fine line between amusement and revulsion. And Kätšyri thinks that the living statues who grace many of our public spaces might be extracting entertainment value from the perceptual mismatch elicited when an apparently inanimate object moves.

The unbearable likeness of beings

The uncanny valley says make your avatar or robot more and more human-like in appearance and it becomes more appealing – but go too far and our reaction sinks into a trough of unease



Laura Spinney is a writer based in Paris, France