

muse


NOVEMBER/DECEMBER 2017

ARTIFICIAL INTELLIGENCE



From Cricket Media

FEATURES



26

Deep in the Uncanny Valley

Why do almost-human robots give us the creeps?

by Kristina Lyn Heitkamp

10

The Mind in the Machine

The making of artificial intelligence

by Kathryn Hulick

18

Testing Thinking

How closely does AI mimic human intelligence?

by Susan Berk Koch

20

Calling the Bluff

A computer bested professional poker players.

by Elizabeth Pagel-Hogan

32

Doctor Robot-o

Will the doctors and nurses of the future be learning machines?


by Lela Nargi

40

How to Deal with a (Seemingly) Racist Computer

Notice blind spots.

by Evelyn Lamb



DEEP IN THE UNCANNY VALLEY

WHAT IS IT ABOUT ALMOST-HUMAN ROBOTS THAT CREEPS US OUT?



In a study coauthored by Maya Mathur, these 80 robot faces were arranged in order from most mechanical to most human.

Have you ever taken a selfie that looks like a doe-eyed deer or a perky-eared dog? The phone app Snapchat's Filters feature uses artificial intelligence (AI) to map your face and track its movements, allowing you to add effects to a selfie. Virtual personal assistants, like Apple's Siri or Amazon's Alexa, answer questions or help out with small tasks, such as playing a song. For most people, interacting with these types of computer programs feels pretty normal. But when it comes to AI robots, it's a whole different game. Humanoids—or robots that act, talk, and look like humans—create complex problems.

People are really picky about what kind of humanoid they will tolerate, much less like. But this sensitivity is nothing new.

DOG AND ROBOT SHOW


Eighteenth-century Swiss watchmaker Pierre Jaquet-Droz had a penchant for crafting elaborate watches and mechanical devices. He also loved to tinker with automatons. In 1759, he presented his automaton of a dog guarding a basket of apples to the court of Spain and its King Ferdinand VI. But when the king helped himself

to an apple, the story goes, the mechanical hound lurched forward and produced such a bark that it incited a real dog in the room to go ballistic. Several courtiers fled the room, fearful that witchcraft was afoot, or that the machine was possessed.


But Jaquet-Droz's best-known robots were three likable human-

A machine that moves automatically, or automaton, captured attention at the court of King Ferdinand VI.





Likability



Mechanical

Human-Like

like figures called the Writer, the Draftsman, and the Musician. Each worked with a simple hand crank. The Writer, an automaton in the shape of a young, barefoot boy, was most popular. The Writer would dip his quill in an inkwell, shake the pen, and then scribe any desired text of up to 40 characters while his eyes followed each written word.

EDISON'S LITTLE MONSTERS

Best known for inventing the incandescent light bulb, American inventor Thomas Edison also tinkered with automatons. In 1890, Edison introduced his talking dolls.

This toy was not a typical cuddle-in-your-arms baby doll. Edison's talking dolls stood almost two feet (61 cm) tall and weighed a heavy four pounds (2 kg).

Customers rejected Thomas Edison's talking dolls.



The metal body held a miniature phonograph. When someone turned the hand crank in the doll's back, she recited a variety of rhymes including "Twinkle, Twinkle Little Star," "Hickory, Dickory, Dock," and "Jack and Jill." Despite Edison's success with other inventions, his dolls were a major bust. They were too expensive, and the voices were strange and hard to understand. The dolls were only in production for about a month. After the disappointment, Edison renamed the dolls his "little monsters." Years later, designers are still decoding why some automatons give people the heebie-jeebies.

ENTER THE UNCANNY VALLEY

When Japanese roboticist Masahiro Mori pondered the future of humanoids, he envisioned robots that looked like the realistic wax figures he had seen in museums. And this idea made him shudder. In 1970, Mori published an article exploring that weirdness. His "uncanny valley" theory plots a



graph outlining people's reactions to different types of automatons, including realistic, wax-doll-like robots. Mori hypothesized that, in general, as a robot's resemblance to humans increases, people's attraction increases. We like robots that remind us of ourselves. For example, we



relate more to *Star Wars*' C-3PO than to an industrial robot that builds cars, or even a stuffed Furby. But as a robot starts to look *almost* human, people begin to have negative feelings toward it. This gap between nearly human and totally human is the uncanny valley. A humanoid's rubbery skin, blank stare, or mismatched movements can appear monstrous. Our response becomes positive again when the robot moves beyond the *almost* human stage and becomes a perfect match for humans. But figuring out the sweet spot has been tricky.

LOCATING THE UNCANNY

A biostatistician at Stanford University's Quantitative Sciences Unit, Maya Mathur grew up reading *Muse* magazine. But nowadays she is trying to detect where the eerie meets the enjoyable in the uncanny valley. Mathur wonders how humans will navigate a social world with robot partners. "I think we will start seeing robots in many everyday places—you may soon find yourself walking around a museum asking questions to a robot museum guide," Mathur says. Plus the medical community is interested in using robots to teach and build social skills with kids who have autism. Robots may soon help us in many other ways as well. But for a robot to be helpful, it can't be scary or untrustworthy. "I would worry that such robots could easily backfire if they provoke uncanny reactions," she says.




Mathur's research attempted to determine the likability and trustworthiness of a variety of robot images, ranging from industrial robots to humanoids. Using 80 still images of real-world robot faces, she asked participants to rate how human each face looked and how enjoyable—or creepy—it might be to interact with the robot in an everyday situation. The images represented a range of design choices

and elements that creators use when building a robot. The results showed that as robot faces progressed from a mechanical design to a human-like design, likability increased—but only up to a point. The sample images followed Mori's curve of eeriness. "I was amazed that the uncanny valley effect was so strong across so many different types of robot faces," she says. Mathur also noticed that when a robot's face seemed to appear happy, it was considered more likable than its negative or emotionless counterparts.

Do you relate more to C-3PO (left) or to a Furby (right)?





A TRUSTWORTHY FACE

Next, Mathur wanted to test how trustworthy humans found the robots. She says it's critical to make sure the humanoids providing important services earn our trust. We must interact positively with a search-and-rescue android or a robotic surgeon. "Feeling mistrustful of and repelled by an android museum guide is one thing, but is quite another if the robot is supposed to be saving your life," she says.

The researchers avoided asking direct questions like "Do you trust this face?" Instead, participants played a game. They engaged in a simple wagering activity, or "investment game." Test subjects were told that money they wagered, between \$0 and \$100, would be given to the robot labs. A pictured robot would decide how the money was distributed. Players reported how much money they would trust to a given robot. The results, again, followed the uncanny-valley pattern. A mechanical-looking robot, including one that had a face like WALL•E but with a large red mouth and big ears, received around \$40, while a creepy-looking,


near-human robot received a wager of approximately \$34. The wager increased to about \$44 when the robot face became identical to a human face. Plus, just like in the likability test, the robots that looked happy were also trusted more than those that showed other emotions. Previous research has shown that people feel a strange eeriness around humanlike objects that display no emotion.

OVERCOMING THE UNCANNY


Paul Brenner grew up loving science fiction stories of robots. He knew as a kid that someday he wanted to build them. Today, he works in the Bristol Robotics Laboratory at the University of the West of England. His work focuses on social human-robot interaction. In 2016, his team tested whether showing emotion reduces the uncanniness of human-like robots. Participants sat in a room with a robot head. The head, named Jules, is the creation of Hanson Robotics, a company that produces very expressive humanoids. Each test

subject was asked to think of a time when he or she was successful, like scoring a high grade on a test, and a time when she or he had failed, like losing a soccer game. Then participants told Jules about each experience. After the story was over, an experimenter in the other room tapped a keyboard control to make Jules respond with either a happy or a sad face. When Jules reacted with a happy face to the story of success, and when it acknowledged the tale of failure with a sad face, its likability score went way up and its uncanny factor dropped significantly. "It was very surprising that adding emotions to a creepy, humanlike head robot impacted the feeling of unease," Brenner says.

While showing emotion helped ease spookiness in this study, Brenner believes the robots that we interact with socially shouldn't be as humanlike as Jules. He says that it's more important to go around the uncanny valley than attempting to overcome it. "While it is useful for robots to be



PROGRAMMING A COMPLEX SHADOW ALGORITHM INTO MY HOLOGRAPHIC DISPLAY REALLY HELPED BRING ME OUT OF THE UNCANNY VALLEY.



Robot designers at the Bristol Robotics Laboratory built RoboThespian to interact with humans in public spaces.

somewhat humanoid, making them appear too human-like is doomed to fail. A robot that is identifiably a machine like WALL•E or C-3PO can be forgiven for making some errors, as people expect this. A robot like Jules gives people too-high expectations of perfectly human-like behavior,” Brenner says. He suggests that robots should be human *enough* so that people can easily figure out how to understand them and interact with them, but not so human-like that people feel tricked. “People don’t like to be fooled, so don’t try and fool them,” he says.

DRAWING THE LINE

As humanoids become familiar figures in our lives, human-robot relationships are increasingly important. Design decisions can make or break the interaction with our mechanical counterparts. But will we ever be completely comfortable with human-like robots? Robotician Masahiro Mori said that perhaps it’s better not to bridge the uncanny valley, but instead stop the design right before it gets creepy. Other roboticists want to create perfectly human-like robots with unique personalities, forging a strong bond between human and humanoid. But what if it becomes challenging to see the difference between a living human and a humanoid? Drawing the line between robot and human may become difficult. But as robot technology evolves, so will our relationships with humanoids.


Kristina Lyn Heitkamp is a children’s book author and science journalist. One of her favorite books is *Tomorrow’s Eve* by French author Auguste Villiers de l’Isle-Adam. The 1886 science fiction novel tells the story of a fictionalized Thomas Edison and his female humanoid creation.

PRINCESS FIONA OR PRINCESS FREAKY?

The uncanny valley affects more than just mechanical robots. Many animated video games and movies are also plagued with accidental creepiness. Hyper-realistic computer-generated imagery (CGI) characters may look human, but their facial expressions can be weird and mismatched. Or sometimes a game’s world is so dynamic that the stiff and dead-eyed characters look out of place and awkward, like in early versions of *Tomb Raider*. “As CGI has gotten better, people expect artificial agents that look human-like to be close to perfect,” Brenner says.

Some CGI films have flopped, driving animators bonkers. The 2004 animated movie *The Polar Express* is notorious for the zombie-like appearance of characters including the train’s conductor, voiced by actor Tom Hanks. Many computer-animation film studios have learned to steer clear of creating human characters that might dip into the strange and bizarre. During a test screening of the 2001 film *Shrek*, kids in the audience began crying when the human version of Princess Fiona appeared on screen. She was realistic in a way that felt uncanny. Luckily, animators caught the mistake and changed her appearance by making her features more cartoony. One film deliberately used the uncanny to induce fear. The 2009 3D stop-motion film *Coraline* wipes out human emotion in the scary characters of “Other Mother” and “Other Father” with blank button eyes and stiff movement. Animators are still searching for a successful bridge between cartoons and the real world.

—Kristina Lyn Heitkamp



VVVVAAAALLEE...

