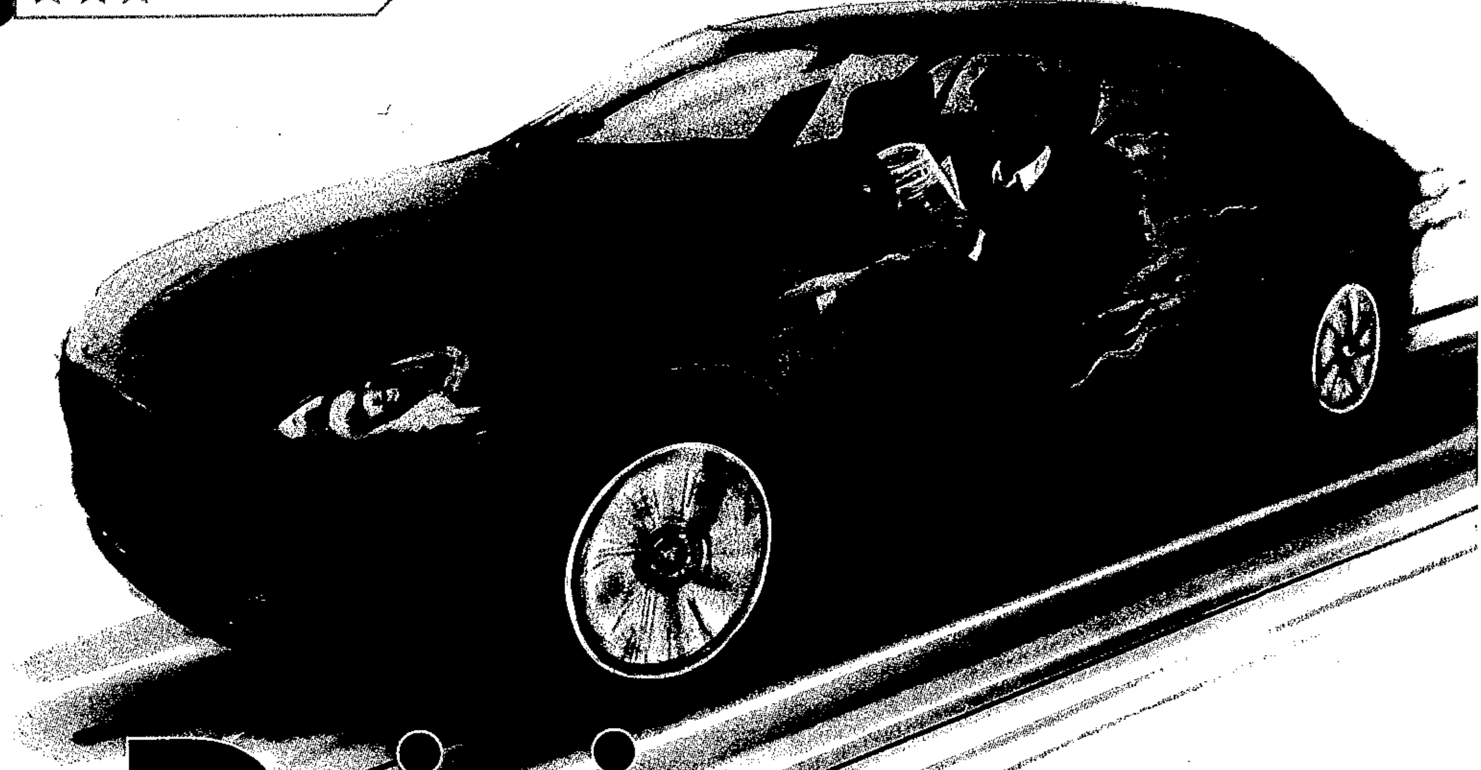

Driving on autopilot

Auto engineers face a paradox: how to design an autonomous vehicle that feels safe but reminds us that no driver, human or artificial, is perfect. Simon Parkin reports on learning to trust a self-driving car

On a clear morning in early May, Brian Lathrop, a senior engineer for Volkswagen's Electronics Research Laboratory, was in the driver's seat of a Tesla Model S as it travelled along a stretch of road near Blacksburg, Virginia, when the car began to drift from its lane. Lathrop had his hands on the wheel, but was not in control of the vehicle. The Tesla was in Autopilot Mode, a highly evolved version of cruise control that, via an array of sensors, allows the car to change lanes, steer through corners, and match the lurching of traffic unaided. As the vehicle lost track of the road markings, he shook the wheel to disengage Autopilot. 'If I hadn't been aware of what was happening, it could have been a completely different outcome,' Lathrop told me recently.

The same week, six hundred miles south of Blacksburg, in Florida, a forty-year-old Tesla driver named Joshua Brown experienced that different outcome. His Model S, driving on Autopilot along Route 27, crunched into the side of an eighteen-wheeler, passing beneath the vehicle's trailer, which sheared off the Tesla's roof and windshield. Brown was killed. The crash is the subject of an ongoing inquiry, but the company was quick to issue an explanation. In a blog post published the day that the accident was first announced, Tesla stated, 'Neither Autopilot nor the driver noticed the white side of the tractor trailer against a brightly lit sky, so the brake was not applied.' The company further noted the 'extremely rare circumstances of the impact'.

Safety record

As Tesla pointed out in its blog post, Brown's is the only death so far in more than a hundred and thirty million miles of Autopilot driving. Google's fleet of similar vehicles has, according to the company, driven more than 1.5 million miles with only one minor collision. Uber, a company for which self-driving taxis may become the full and final act of putting cabbies out of work, has a test car on the road in Pittsburgh that it hopes will make 'transportation as reliable as running water'. Indeed, so far, autonomous vehicles have had an exemplary safety record. Tesla's data demonstrates a slight improvement over humans but it is a flawed comparison. Autopilot is designed for use only on freeways, where human drivers, too, have far fewer accidents. And even if travelling by autonomous vehicle is shown to be statistically safer, car designers will need to find ways to reassure people beyond mere numbers.

'The real hurdle to the widespread adoption of autonomous vehicles is psychology,' Chris Rockwell, the CEO of Lextant, a research consultancy that focuses on user experience, told me. 'People will forgive other humans much more quickly than they will technologies when they fail.' At Volkswagen, Lathrop is currently working on the problem with Traffic Jam Pilot, which is expected to feature in the next Audi D-Class. The system can control the car and issue a warning should a human be seen by its cameras to fall asleep at the wheel, a provision similar to Tesla's requirement that drivers keep their hands on the wheel, even when in Autopilot.

Trust

'There are three key ways to make the occupants of a self-driving car feel safe,' Lathrop said. 'It must be clear when the vehicle is operating in autonomous mode. Occupants must know that the car is sensing its environment, other vehicles, and pedestrians and so on. Finally, the vehicle must prime people before it makes a manoeuvre.' There's nothing more disconcerting for passengers than when a driver makes abrupt lane changes or swerves. According to a paper published in 2014 in the *Journal of Experimental Social Psychology*, the more humanlike the car's alert features (name, voice, gender) the more people trust it to operate competently, as a human driver would.

The problem of trust faces outward as well as inward. How do pedestrians and other drivers distinguish a vehicle that is driving autonomously from one that is not? Recently Google was awarded a patent for an adhesive hood, designed to stick a human to the front of a self-driving car, preventing secondary injuries caused by tumbling into the windshield or rebounding onto the asphalt. Volkswagen's solution to instilling pedestrian trust is rather more mundane. The company has tested an autonomous Audi A7 that features a strip of LEDs facing out of the front windshield. The lights blink and follow pedestrians at a crosswalk to signal that the car sees them, the equivalent, perhaps, of a friendly wave of the hand.

Staying alert

A decade ago, according to Lathrop, the company ran a series of internal studies in which it put people in the driver's seat of a car that they were informed was fully autonomous. Behind them, behind a curtain, sat a driver, who controlled the car using a camera feed of the road ahead, as if playing a video game. 'We found that people get comfortable very quickly, almost too quickly, in fact, in letting the car drive itself,' Lathrop said. Freed from the activity of driving, most people experience what researchers refer to as passive fatigue, a state in which awareness is dulled. It can set in after as little as ten minutes. While wearied by inactivity, a car's occupants typically look for distractions.

Frank Baressi, the sixty-two-year-old driver of the truck that killed Brown, claims that when he approached the wrecked Tesla he heard one of the *Harry Potter* films playing inside the car. Investigators found both a portable DVD player and a laptop inside. Tesla strenuously warns consumers to pay attention while their car is in autonomous mode, but the warning may not be strong enough. Lathrop is working with his colleagues at Virginia Tech on using drivers' phones or tablets to pass along alerts while they're at the wheel, to make them harder to ignore.

Paradox

But whether or not Brown was distracted at the time of the collision, Lathrop said that, as autonomous systems improve and trust in them increases, the temptation for occupants to do other things will grow stronger. 'We are not naïve,' he said. 'But ultimately the operator of the vehicle is responsible for having some degree of situational awareness. When it comes to autonomous cars, it's a system. It's a machine. It's not aware of everything. It's simply sensing its environment and responding as it has been trained.'

This is the paradox facing auto engineers: how to design self-driving cars that feel trustworthy while simultaneously reminding their occupants that, no driver, human or artificial, is perfect. How to free drivers from the burden of driving, while burdening them with the worry that, at any moment, they will need to take back control. <<

FRANKENSTEIN 2.0

Two hundred years after Mary Shelley wrote Frankenstein, we look at how today's scientists can create life in the lab

By Duncan Geere

In June 1816, a monster was born. Mary Shelley (then Mary Godwin) was holidaying along the banks of Lake Geneva with Lord Byron and her lover, Percy Shelley. It was a cold, wet summer, so, stuck indoors, Byron challenged everyone in the group to write a ghost story. A few days later, Mary Shelley began working on what would become *Frankenstein*. Shelley's story was undoubtedly influenced by the science of the day, but what would have inspired her if she were alive today? Regenerative medicine and biotechnology are advancing at a breakneck pace, and the idea of creating life in the lab is looking less and less implausible. Here we look at the science that's making Mary Shelley's vision a reality 200 years on.

The hardest part of building a creature from scratch, after the brain, is the head and face. While head transplant experiments have been carried out on animals for more than a century, all attempts have ended in paralysis and the animal's death. The most influential researcher in this area was scientist Vladimir Demikhov, who experimented with dog head transplants in the Soviet Union in the 1950s. He was unsuccessful, but his other experiments in transplanting organs between animals significantly advanced the field, including the use of immunosuppressants to reduce the risk of a body rejecting the transplanted organ. His work led directly to the first human heart transplant in 1967.

SWAP HEADS

In 1970, a team of researchers led by US neurosurgeon Robert J. White attempted to transplant the head of a monkey onto the body of another. The procedure was a partial success, with the animal surviving for some time after the operation and reportedly able to sense the world around it, but the public greeted the news with widespread disapproval.

More recently, a new figure has appeared on the scene. Sergio Canavero, an Italian neurosurgeon, has attracted widespread media attention over claims that he'll perform the first successful human head transplant in 2017, with some calling him the 'real-life Frankenstein'. Canavero already has a patient, a 30-year-old Russian named Valery Spiridonov with spinal muscular atrophy. The surgeon's plan involves slicing off Spiridonov's head using a clean, fast procedure, and then connecting it to the donor body's spinal cord with a polyethylene glycol 'glue'. But there is much doubt, not only over whether he'll succeed, but also whether he'll even be able to attempt it in the face of financial and ethical constraints.

MAKE A FACE AND BODY

If heads are too hard, then faces are easier. In November 2015, Patrick Hardison, a firefighter who had been horribly burned in an accident, was given the face of a brain-dead man during a 26-hour-long operation at New York University's Langone Medical Centre. The surgery came almost exactly 10 years after the first partial facial transplant in 2005 and was described as 'a critically important contribution to the advancement of science and medicine' by the medical centre's dean Robert Grossman. Recipients of facial transplants must take immunosuppressant drugs for the rest of their lives, and are at greater risk of suffering from infections and cancer.

Once the brain and head are constructed, the rest of the body is easier to build. Bones are probably the simplest of all – we've been fixing bones with pins, rods, and screws since the middle of the 19th century. Frankenstein sourced his bones from 'charnel houses', but he would have loved 3D printers, which can form perfect replica bones of titanium. In 2012, an 83-year-old woman was given a 3D-printed lower jaw that took just a few hours to print and install. The patient was able to speak shortly after waking up from the anaesthetic. A similar titanium ribcage was installed in a Spanish patient in 2015. Organs are a little harder, but we're getting there. Bioprinters have been able to create human tissue for a while. Artificial hearts, kidneys, and livers have all been printed.

STAYING ALIVE

The difficulty with 3D-printed tissue is keeping it alive, as this requires tiny blood vessels. In 2014, scientists in Australia and the US took the first steps towards integrating artificial blood vessels into tissue, and only recently researchers at North Carolina's Wake Forest University published the details of a 3D printer that could create everything: organs, tissues, and bones that could all be implanted into humans. The field is moving fast, but most experts warn against optimism, saying it'll likely be decades before we see the tech becoming common.

That just leaves skin, which is surprisingly easy to print, thanks to its layered structure. In 2015, L'Oréal announced that it was teaming up with bioengineering start-up Organovo to 3D-print human skin. The companies said the skin would be used in product tests, though some medics have suggested it might have more value in burns units and trauma centres.

THE FINAL CHALLENGE

Now we're back to where we began, the 'spark' of life itself, something which we're not much closer to understanding today than during Mary Shelley's time. Dr Frankenstein famously discovered the secret of life: 'I became myself capable of bestowing animation upon lifeless matter.' But no real-world scientists have come close. So animating our bundle of 3D-printed organs, transplanted head, and digital brain is no easy task. Life is infinitely more complicated than a chemical reaction, or software algorithm; that's why we can only approximate it in the lab, not replicate it entirely.

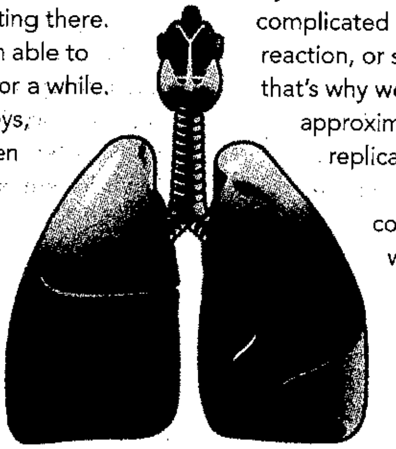

But let's say we could wave a magic wand to create that vital spark. How would it interact with the body we've created for it? Mind-controlled prostheses are already a reality; researchers at Johns Hopkins University recently announced a prosthetic arm whose individual fingers can be controlled by the brain. To configure it, electrodes were implanted over the part of the brain that controls hand and arm movements, with researchers tracking the locations that emitted an electric pulse when the subject moved their fingers. These signals could then be used to trigger movements in the prosthetic hand in the same way.

SIMULATE SENSES

Currently, there are prosthetics that can deliver the sensation of touch. In 2013, University of California biologists connected up the brains of monkeys to an artificial fingertip equipped with sensors, using a similar brain-location-mapping process. They found that the monkeys responded the same way to 'feeling' in the artificial finger as in their real fingers. Fredrik Winquist at Linköping University in Sweden has built an electronic tongue that can differentiate between tastes, while Massachusetts-based firm C2Sense has created a similar device for smell. Combined with cameras and microphones for sight and sound respectively, that's pretty much all the body's major senses covered.

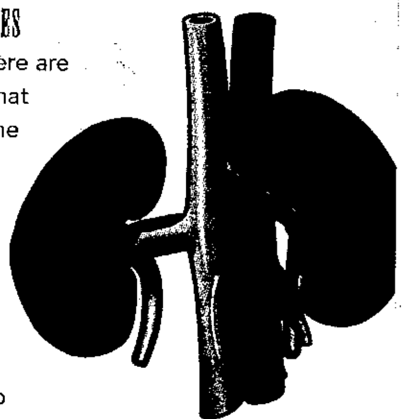

Much work remains to be done, not least teaching a digital brain how to cope with these inputs and process them into actions. It's easy for us to quickly withdraw our hand when it touches a hot surface, but harder for a computer to perform all the calculations at a speed that avoids damage, while also processing continuous input from the rest of its body. A fully functioning artificial being might be some way off yet, but it's surely only a matter of time. <<