How AI changed organ donation in the US

qz qz.com/1383083/how-ai-changed-organ-donation-in-the-us

September 10, 2018

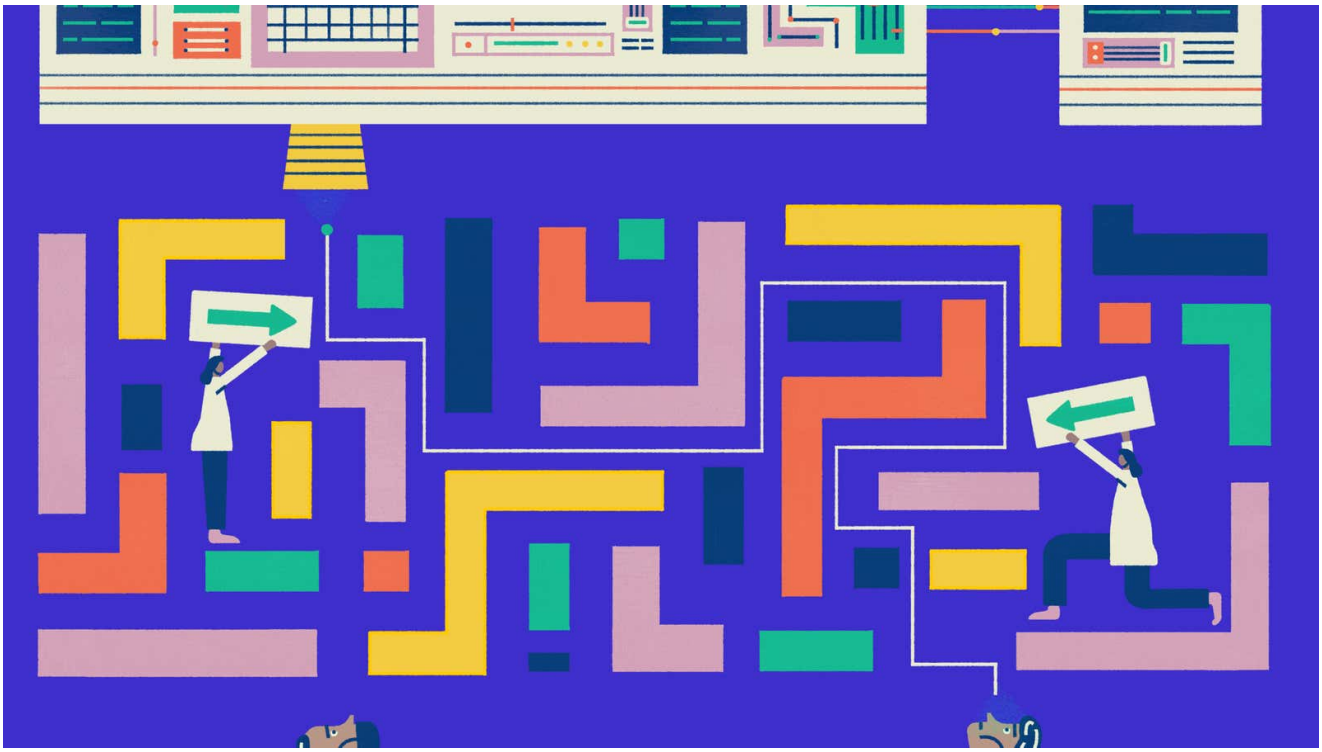

There used to be only three ways off of a kidney transplant waiting list. The first was to find a healthy person from within one's own pool of friends and family, who perfectly matched both the recipient's blood and tissue types, and possessed a spare kidney he or she was willing to part with.

The second was to wait for the unexpected death of a stranger who was a suitable physical match and happened to have the organ-donor box checked on their driver's license.

The third was to die.

But then it occurred to doctors: given enough kidney patients, and enough healthy, willing donors, they could form a pool big enough to facilitate far more matches than the one-to-one system of the past. As long as patients could procure a donor—any donor, even one that wasn't a fit with the patient themselves—they could get a matching kidney.

At first, this required doctors to spend brain-searing hours poring over the details of blood types and tissue variations in patients' and potential donors' charts. Then computer scientists and economists got involved. They built algorithms that performed these complicated

matches more elegantly than human brains ever could. Now, thanks to artificial intelligence, a person stepping forward to donate a kidney to a loved one—or to a perfect stranger—can set off a chain that saves dozens of lives.

Paired kidney donation is one of the great success stories of artificial intelligence. It doesn't eliminate jobs or scrub the human touch from medical care. It takes an incredibly complex problem and solves it faster and with fewer errors than humans can, and as a result saves more lives. Since the first paired kidney exchange surgeries took place in 2000, nearly 6,000 people have received kidney transplants from paired exchanges identified by algorithms. Today roughly one in eight transplant recipients who receive a kidney from a living donor are matched with that person through paired exchange.

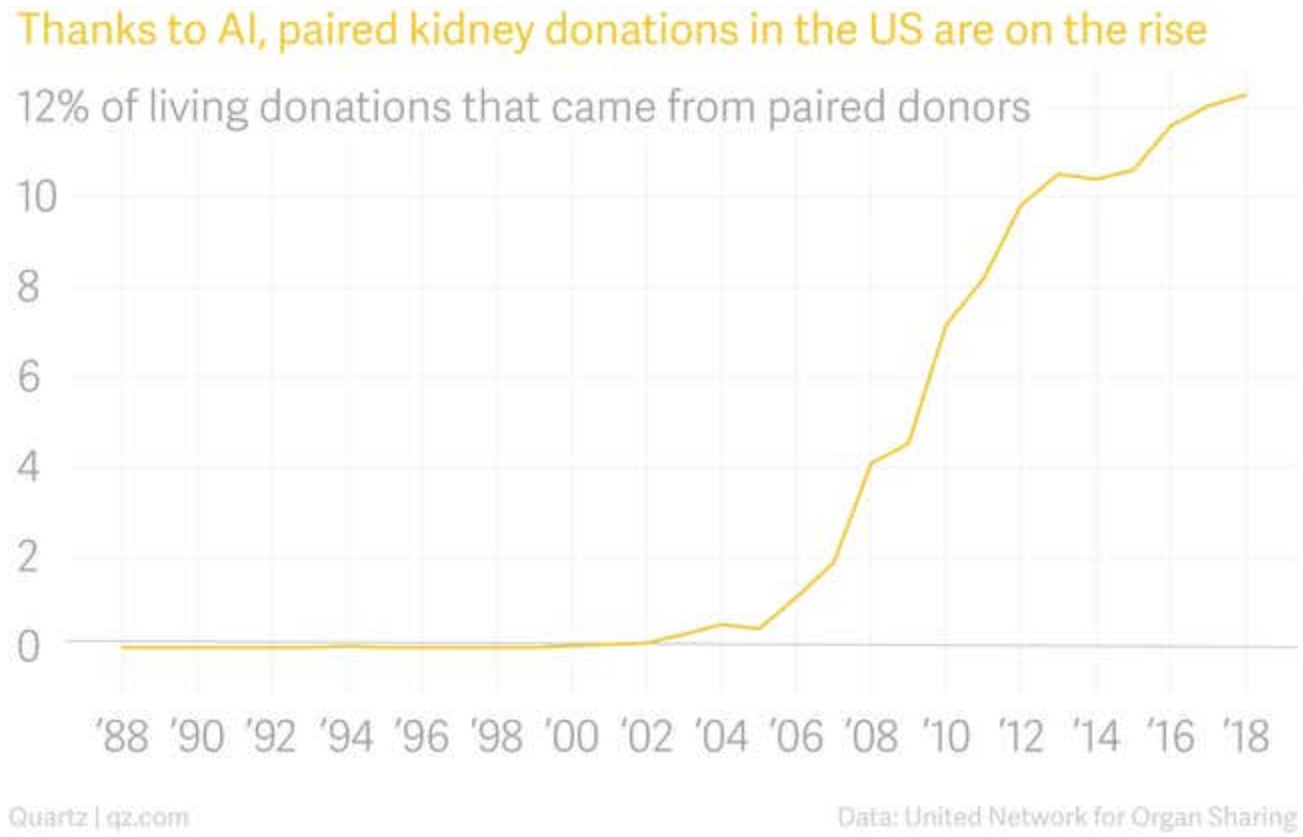

At the same time, paired kidney exchange is also a perfect example of AI's limitations. A computer can only do what a human can teach it, and we can't teach what we don't understand. In the decades since medicine learned how to replace a failing kidney with a donated one, we are still struggling with the problem of how to distribute the precious few kidneys available in a way that feels fair and satisfactory to everyone, and doesn't result in undesirable, unintended consequences. AI can identify potential donors and recipients who are biologically suited for one another; in the future, it may even be able to weigh the moral factors that determine who gets a transplant first. But first, we humans have to agree on what those should be.

The kidneys act as the body's filters. For people with kidney failure, dialysis essentially replicates the organs' function externally, removing a patient's unfiltered blood over a period of hours and pumping it back in. The invention of dialysis in the middle of the 20th century turned a disease that was once a death sentence into a chronic but manageable condition.

The first outpatient dialysis center, the Seattle Artificial Kidney Center, opened in January 1962. Because each patient needed to be hooked to the machine for two 12-hour sessions per week, the center could accept only 10 patients, of the roughly 2,000 end-stage renal patients in the US eligible for dialysis at that time, for its first two years.

The "God committee"…consisted of a lawyer, minister, banker, housewife, state government official, labor leader, and surgeon.
To distribute the coveted spots, the center organized an independent committee of seven citizens. The first Admissions and Policies Committee—or the "God committee," as it was later referred to in the press—consisted of a lawyer, minister, banker, housewife, state government official, labor leader, and surgeon. Doctors made some decisions for them: no patients over 45 were eligible, nor were children, who doctors worried might be traumatized by the procedure. Apart from that, the committee was left to its own devices to choose which of the roughly one in four eligible patients who applied for a spot to save.

The anonymous committee members looked at applicants' age and gender, and whether they were married or had children, according to a 1962 Life magazine article. They looked at their emotional stability. They looked at how much they earned, and how much of those earnings they'd saved; their education level, their job, their past deeds, and future potential. They decided that since the research to develop the technology had been done at institutions funded by the state of Washington, only Washington residents would be considered, as it had been their taxes that paid for the treatment.

The article described the committee members as "high-minded, good-hearted citizens" facing a nearly impossible task. They were also human and, as humans, subject to conscious and unconscious biases. They were making decisions based on case histories written by doctors with their own conscious and unconscious biases. Life magazine reporter Shana Alexander transcribed a section of their discussion (ellipses in original):

> HOUSEWIFE: If we are still looking for the men with the highest potential of service to society, then I think we must consider that the chemist and the accountant have the finest educational backgrounds of all five candidates.
>
> SURGEON: How do the rest of you feel about Number Three—the small businessman with three children? I am impressed that his doctor took special pains to mention that this man is active in church work. This is an indication to me of character and moral strength….
>
> HOUSEWIFE: Which certainly would help him conform to the demands of the treatment….
>
> LAWYER: It would also help him endure a lingering death….

> MINISTER: Perhaps one man is more active in church work than another because he belongs to a more active church.
>
> BANKER: We could rule out the chemist and the accountant on economic grounds. Both do have a substantial net worth….
>
> LAWYER: Both these men have made provisions so that their deaths will not force their families to become a burden on society….
>
> STATE OFFICIAL: But that would seem to be placing a penalty on the very people who perhaps have been most provident….

> MINISTER: And both these families have three children too.
>
> LABOR LEADER: For the children's sake, we've got to reckon with the surviving parent's opportunity to remarry, and a woman with three children has a better chance to find a new husband than a very young widow with six children.
>
> SURGEON: How can we possibly be sure of that?

The seven committee members were a cross-section of US Pacific Northwest upper-middle-class respectability, and their discussions (at least those published) hint at a favorable regard for patients who seemed to share their status and values. Race is never mentioned in the Life article, but it is hard to imagine that, in a society still deeply segregated (to give just one example, 98% of Washington state employees at the time were white), people of color had anything remotely close to a representative seat at the table.

"As a human being, do I have that right? I don't really think I do."
In interviews, the committee members sounded overwhelmed and troubled by the power they'd been handed. What, one member mused, if they were approached by an extremely wealthy patient who would agree to fund the treatment of all patients who followed him in exchange for access to the machine tomorrow? Does an artist have more future value than a

pastor? “I do have reservations about the moral aspects, the propriety of choosing A and not B, for whatever reason,” the banker said. “As a human being, do I have that right? I don’t really think I do.”

Their deliberations foreshadowed a problem that would persist long after that primitive kidney treatment machine was surpassed by more sophisticated technology. Even when a group of people is unanimously committed to doing the best thing possible, for the greatest number of people, what, exactly, does the best thing look like?

While the God committee issued its wrenching decisions, a parallel effort to treat kidney patients was underway elsewhere.

Doctors performed the first successful kidney transplant in 1954, at what is now Brigham and Women’s Hospital in Boston, surgically removing an organ from a 23-year-old man and giving it to his identical twin brother, who lived another eight years.

By the early 1960s, developments in blood and tissue typing made it easier for doctors to identify successful matches between donors and recipients, and improved immunosuppressant drugs dramatically lowered transplant-rejection rates. Today, a kidney from a deceased donor functions in the recipient’s body for another eight to 12 years, while one from a living donor lasts an average of 12 to 20.

In 1972, US president Richard Nixon signed legislation expanding Medicare to cover dialysis for all patients with kidney failure. Dialysis didn’t cure them, but it kept more people alive for longer while they waited and hoped for a kidney transplant. From then on, the number of people needing kidney transplants in the US grew much faster than the number of donor kidneys available.

As of writing, there are 114,554 people awaiting organ transplants in the US, according to the United Network for Organ Sharing. And 94,980 of them—83%—are waiting for a kidney.

The 1968 Uniform Anatomical Gift Act established a standard procedure across the US through which people or their next of kin could authorize donation of their organs after death. But even if everyone in the US was a registered organ donor (the current rate is 54%) there wouldn’t be enough kidneys to meet the need. Less than 2% of people die in a way that makes them suitable organ donors. For organs to remain transplantable after a person’s death, oxygenated blood must be continually pumping through them until they are surgically removed from the body. Deceased donors have typically been declared brain dead and placed on a ventilator before the procurement surgery.

Fortunately, since most humans are born with two functioning kidneys but only need one to survive, living people can donate kidneys. As kidneys from living donors also tend to last longer than those from dead ones, finding a living donor is usually the ideal outcome for a

person in need of a transplant.

While all surgeries carry some risk of complication or death, the vast majority of kidney donors spend just two or three days in the hospital after the laparoscopic procedure, followed by an additional four to six weeks' recovery time. The remaining kidney grows to compensate for the missing one, and donors typically go on to live normal lives.

Petitioning your sibling or cousin for an organ is no small ask, of course. But finding a willing donor is often the easiest part of the process, as many patients and their families are chagrined to learn.

Florida businessman Neil Emmott was diagnosed in 2001 with polycystic kidney disease, a genetic condition that can lead to kidney failure. The news was "unexpected and devastating," his wife Lisa Emmott says. At the time of the diagnosis, Neil was 38, Lisa was 27, and the couple had been married less than a year. By 2016 they had two young daughters, and Neil's condition had deteriorated to the point that his doctors at Johns Hopkins University recommended he research transplant options.

Lisa volunteered to donate immediately. She was healthy and shared his blood type, and so assumed she'd be a viable candidate. But it takes more than a blood and tissue match for an organ donation to proceed. Donors are thoroughly screened for any underlying physical, psychological, or socioeconomic problems that might complicate their donation. Lisa learned that a benign quirk in the shape of her renal arteries—the veins that carry blood to and from the kidneys—disqualified her. Neil's brother stepped up as a substitute, only to be ruled out as well due to minor medical issues. The family was heartbroken.

"Needing an organ is an awful place to be," says Emmott. That's when the family decided to turn to the marketplace.

Neil and Lisa Emmott in Sept. 2017.
Image: Courtesy of Lisa Emmott

There's something creepy about the idea of a "market" for human body parts. They're not commodities: US law unambiguously bars the selling of human organs. But in economic terms, a market is just any place where people who want something find people who can give it to them. A market that doesn't rely on price alone to allocate resources is known as a matching market. A dating pool is one type of matching market (assuming no money is being exchanged for companionship); so is the pool of people who want kidneys and people willing to give one up.

Markets work best when they're "thick," or have a lot of participants. For the first several decades of kidney transplantation, sick people and their prospective donors were limited to the very thin market of just themselves. Failed matches often proved a death sentence for the seriously ill.

But what if the market could be made thicker?

Image: Quartz

This idea was first floated in a 1986 paper by the German-born, US-based surgeon Felix Rapaport, who theorized that it could be possible to transplant kidneys across two willing donor-recipient pairs: Patient A receives a kidney from Donor B, and in exchange, Donor A gives a kidney to Patient B.

In 1991, doctors in Seoul, South Korea, under the direction of nephrologist Kiil Park, performed the first paired kidney transplants between two donor-patient pairs. Four years later, the world's first paired kidney donation program opened at the Yonsei University College of Medicine in Seoul. Potential donors and recipients were entered into a database and then manually paired by doctors through hours of painstaking analysis. In 1999, Switzerland became the next country to establish a paired kidney exchange, matching up two married couples that each had one spouse with end-stage renal disease and one spouse willing to donate a kidney.

One night in 2000, tired of delivering the heartbreaking news to patients and their loved ones that no suitable kidney could be found, a US nephrologist named Michael Rees lugged home several crates of files and spent the next few hours scrutinizing blood, antibody, and tissue data, and comparing patient charts. The work was mentally grueling. Eventually, he realized he had no viable matches—but also, that if the pool were bigger, pairs could be made. Working with his father Alan Rees, a computer scientist, Michael Rees created a simple computer program that did the work of pairing up donors and recipients, introducing AI to the matching process.

Around the same time, Alvin Roth, a professor of economics at Harvard University, was also tinkering with kidney-matching solutions. Roth specialized in market design, focusing on how markets could be adjusted to fix supply-and-demand imbalances. He'd previously designed algorithms to match new doctors to residency programs, and New York City grade-school students to high schools. Now he turned his attention to kidneys.

Roth and his colleagues Utku Unver and Tayfun Sonmez designed an algorithm that reviewed and analyzed the data profiles of potential donors and recipients. It identified "cycles" of donor-recipient pairs, and "chains," in which a person who chooses to altruistically donate a kidney to anyone who needs one (and such donors exist) sets off a sequence of donations from among the pool of potential donors and recipients registered

with a hospital or kidney exchange. For example: Patient B receives a kidney from Altruistic Donor A, after which Donor B gratefully donates a kidney to Patient C. If Patient C has a willing donor, there's no real limit to how long the chain can grow. Unlike cycles, chains could move forward indefinitely, without having to go back and close the loop by finding a kidney for the original donor's partnered recipient.

Image: Quartz

Roth, Unver, and Sonmez felt like they were on to something big. In 2003 they posted a paper outlining their work online and sent it to nephrologists across the US. After feedback from a Harvard surgeon named Frank Delmonico, the team adjusted their algorithms and published a new paper whose concepts helped establish the New England Program for Kidney Exchange. The exchange matched donors and recipients across the 14 kidney transplant centers in the region.

At first, surgeons insisted all the surgeries in a given cycle or chain take place simultaneously, so that no donor could back out at the last minute. This limited the number of patients in a cycle or chain, as there were only so many beds and so many surgeons any hospital could spare at once. This, the economists and several other doctors argued, was unnecessarily restrictive. There were no biological barriers: unlike hearts or lungs, which must be transplanted within four to six hours of leaving a donor's body, a kidney can be safely preserved for 24 to 36 hours before finding its new home. As for the possibility of a weak link in the donor chain, the economists argued that in a chain starting with a donor willing to give their kidney to anyone, no recipient would be left stranded if a donor got cold feet, as doctors could find a replacement from within the pool of registered donors.

Rees, the nephrologist who built the first matching algorithm, proved it could be done. After a 28-year-old came to him with the offer give a kidney to a stranger in need, Rees organized a chain of kidney donations that saved the lives of 10 patients across five states over an eight-month period.

Today, multiple US hospitals run their own paired kidney donation programs. There are also three larger US exchanges that organize kidney chains across hospitals: the United Network for Organ Sharing, the National Kidney Registry, and the Alliance for Paired Kidney Donation. National exchanges are in place in the UK, Canada, and the Netherlands, and paired donations have taken place in hospitals from India to South Africa. Researchers have also theorized that similar exchanges are possible for lung and partial liver transplants, though no system for such swaps is yet in place.

In 2012, Roth won the Nobel Prize for his work on market design. He brought Rees with him to the ceremony. By then, 2,000 people in the US had received transplant as a result of the system they helped create. Thousands more have since.

Neil Emmott ended up being one of eight people in a 2017 kidney chain that started when two family friends came forward to donate in his honor. On Aug. 13, 2018, a woman in Alabama became the 100th person to receive a kidney in a nationwide chain that started in 2013.

Left, Allison Malouf holding a photo of her kidney shortly after its removal in November 2017. Right, Britani Atkinson showing her nephrectomy scar. Malouf and Atkinson were part of the paired-donation process that enabled Neil Emmott to get a new kidney.
Image: Alison Frank Photography

Today, when doctors are looking to match kidney donors and recipients, algorithms built by AI researchers trawl the database of registered kidney patients and their partnered donors, and identify matches based on a list of weighted criteria hammered out by a committee of the Organ Procurement and Transplantation Network and the United Network for Organ Sharing, which together oversee US organ transplants.

The algorithms evaluate all the transplants possible among the patient-donor pool at once. Matches are made primarily on biological suitability, with the hardest-to-match patients getting first priority. The technology weighs criteria including the time the recipient has been on the waiting list, his or her age (children get priority), and whether the person who needs a kidney has been a living organ donor in the past (with the reasoning that people who have stepped up to give before should get priority if they one day find themselves in need).

Image: Zack Rosebrugh for Quartz

These algorithms have facilitated thousands of life-saving surgeries. And in the future, it could be possible for an AI not just to make matches using the criteria that humans have decided upon, but to actively participate in that judgment process—to understand human decision making and value systems such that it can make its own tie-breaking judgment calls about which kidneys should go where (a decision that would then be reviewed by human doctors). On this point, the limiting factor isn't so much the technology as the people using it.

The first issue is human anxiety around AI's role in organ allocation. Hospitals and organ exchanges are reluctant to even use the term "artificial intelligence" in conjunction with the matching process. That's in part due to something researchers call the "AI effect," the tendency to apply that label only to futuristic and novel technologies. As Nick Bostrom, the head of Oxford University's Future of Humanity Institute, has put it, "Once something becomes useful enough and common enough, it's not labelled 'AI' anymore."

Given the dearth of public education on what "artificial intelligence" actually means, hospitals and exchanges are wary of patients misconstruing the role algorithms play in identifying potential matches, perhaps fearing conjuring images of robots coldly issuing life-or-death edicts.

Machines currently do not decide which kidneys go where. Humans do that. The algorithms in place today can do the math more reliably and at greater scale than humans can, and implement the judgments humans have already made, but they don't have a contextual understanding of why they are being asked to perform a calculation in the first place.

"AI['s] don't have a general understanding of the world like we do. They don't understand what the data they're processing is about," says Vincent Conitzer, a professor of computer science, ethics, and philosophy at Duke University. "They don't have a concept that this person is suffering. They don't really understand what a person is. Humans have to come in at some point."

Image: Zack Rosebrugh for Quartz

Researchers are now teaching machines to understand these sorts of moral dilemmas from a human perspective. Conitzer, along with Duke colleagues Jana Schaich Borg, Walter Sinnott-Armstrong, and Rachel Freedman, and John Dickerson at the University of Maryland, published a paper this year in which they presented research subjects with hundreds of pairs of hypothetical patient profiles and asked which of each pair should get the one available kidney. Instead of the blood and tissue data that algorithms crunch now, these hypothetical patient profiles listed things like how often the patient drank alcohol, and whether they'd had cancer in the past. The researchers then fed the subjects' choices to an algorithm designed to detect patterns in the human responses, and learn how to choose the "right" kidney recipients in accordance with those patterns. Just like the human subjects, the AI favored younger, healthier patients—an example of a machine making a decision based on what it has learned about human values.

A machine can be taught to make matches in line with our value system, but we don't always understand what our own values ar

A machine can be taught to make matches in line with our value system, but we don't always understand what our own values are, or agree on them as a group. People don't always know what they want to optimize for, and even when they think they do, they often don't understand how to do it in a way that doesn't lead to unintended consequences.

For example, MIT's Moral Machine allows visitors to the lab's website to play a game in which they must choose, in hypothetical case after hypothetical case, which group of car passengers or pedestrians a driverless car should opt to kill when faced with two terrible choices. After clicking through a series of stomach-churning scenarios—*yes, I'd rather the car run over one pregnant woman instead of five homeless adults; no, I wouldn't have it*

swerve to avoid two children if it meant killing five adult passengers—the game reveals the patterns it has identified in your choices, and how your responses compare to those of other players.

Image: Zack Rosebrugh for Quartz

This information can reveal surprising unintended consequences and uncomfortable unacknowledged biases. You might learn, for example, that your decisions disproportionately resulted in the death of more men than women, or that you tend to value adherence to traffic laws above human life far more than the average player of the game.

In the kidney question, take the ostensibly fair-minded principle that kidneys should be given to the people who will likely have the most years of productive life after receiving one. Before a computer can calculate the longevity of potential recipients, scientists have to feed the algorithm data on life expectancy for various populations. But this leads to some problems. Men tend to die before women do. Black Americans die younger than Americans of any other ethnicity. A 65-year-old white woman in the US could expect to live another 20.5 years in 2015, four years longer than a black man of the same age. What started with good intentions ends in systematic racial and gender discrimination. Back to the drawing board we go.

"In economics we talk about impossibility theorems. There are things you might want that are not possible to get," Roth says. "When you're allocating scarce resources, you can't give a kidney to one person without failing to give it to someone else.... Computers will not lift the burden from humans in every respect."

Economist Alvin Roth and nephrologist Michael Rees at the 2012 Nobel Prize ceremony, where Roth received the prize in economics for his work on market design
Image: Courtesy of Alvin Roth

AI didn't cause these ethical dilemmas. Human committees agonize over the fairest way to allocate kidneys; human drivers still have to make terrible emergency decisions behind the wheel. Machines can model in hours or even minutes consequences of human judgments that might have otherwise taken years to spot. That can be helpful in the long run: if a computer model can demonstrate that a particular kidney-allocation policy will disproportionately disadvantage certain groups, for example, doctors can scrap the plan before anyone has actually been harmed. But not everyone is comfortable letting a machine participate in a life-or-death decision.

The Moral Machine teaches players something else that Seattle's God committee learned long ago: having to choose which lives to save, knowing that the decision will result in someone else's death or suffering, feels horrible. This part AI can't help with. People can figure out the

ideal way to assemble a chair and then teach that process to a machine that goes on to assemble thousands of chairs perfectly. There is no perfect way to decide who lives or dies.

This story is one in a series of articles on the impact of artificial intelligence on health care and medicine. Click here to sign up to get alerted when new stories are published.