encourage inclusiveness. These include an updated code of conduct that forbids event sponsors from using sexualized clothing or costumes, a town-hall meeting to discuss the issues, on-site childcare and stickers that help inclusion by flagging up first-time attendees and highlighting the pronoun that people prefer to be referred to by. In addition, more specific meetings for under-represented groups ran alongside the conference than in previous years.

They are small steps down a long road. Most participants easily adopted the name NeurIPS, with only a few accidentally slipping up and mentioning NIPS. There were no offensive T-shirts at conference events and no supply of commemorative coffee mugs — at least when the conference opened — as the name change came too late to get them printed.

Too often, the burden of work involved in increasing diversity and inclusivity falls on those from under-represented groups. NeurIPS is no exception.

One of the organizers of the Black in AI workshop at NeurIPS, Timnit Gebru, a computer-vision researcher at Stanford University, California, spoke for many of the session organizers when she told the diversity town-hall meeting that coordinating the event had reduced the time available for her research. The diversity and inclusion co-chairs, Katherine Heller and Hal Daumé, who have had to walk the fine line between a vocal research community pushing for change and a conference board that has been slow to realize the significance of its actions, also say they have seen considerable disruption to their research. Only Heller, who is at Duke University in Durham, North Carolina, has so far committed to returning to the post next year. These examples underscore the fact that increasing diversity is a job for everyone and it is not sustainable or fair to rely on a small number of volunteers to do this important work.

The challenge should not be underestimated. The organizers of another major AI conference, the International Conference on Learning Representation, announced last month that they would hold their 2020 event in Addis Ababa in a bid to widen the pool of talent that can attend. But this well-meaning initiative is not without problems. Some members of the LGBTQ+ community raised concerns about holding the meeting in a country with anti-homosexuality laws. Organizers assured them of their safety, and that many scientific conferences had been held there before.

"Increasing diversity is a job for everyone, it is not fair to rely on a small number of volunteers."

The Canadian government has come under fire for denying entry or being slow to approve visas for many researchers invited to attend NeurIPS from overseas. More than half of the 200 people who sought visas to attend the Black in AI workshop did not receive them in time, including several who dedicated huge amounts of time to organizing the event.

There have been many high-profile criticisms of AI algorithms that mimic, and so perpetuate, the biases of wider society. That the board of an AI conference — which in October called in a diversity and inclusion consultant to assist it — has taken a stand against such discrimination within its own ranks is a necessary and overdue step. Ensuring changes are deep and lasting will take much more time and effort. Meanwhile, many more institutions and organizations need to follow. ■

How we forget

From pop music to tennis stars, society loses interest according to a mathematical law.

In his enthralling 2009 collection of parables, *Sum: Forty Tales from the Afterlives*, the neuroscientist David Eagleman describes a world in which a person only truly dies when they are forgotten. After their bodies have crumbled and they leave Earth, all deceased must wait in a lobby and are allowed to pass on only after someone says their name for the last time. "The whole place looks like an infinite airport waiting area," Eagleman writes. "But the company is terrific."

Most people leave just as their loved ones arrive — for it was only the loved ones who were still remembering. But the truly famous have to hang around for centuries; some, keen to be off, are with an "aching heart waiting for statues to fall".

Eagleman's tale is an interpretation of what psychologists and social scientists call collective memory. Continued and shared attention to people and events is important because it can help to shape identity — how individuals see themselves as part of a group — and because the choice of what to commemorate, and so remember, influences the structures and priorities of society.

This week in *Nature Human Behaviour*, researchers report a surprising discovery about collective memory: the pattern of its decay follows a mathematical law (C. Candia *et al. Nature Hum. Behav.* http://doi.org/cxq2; 2018). The attention we pay to academic papers, films, pop songs and tennis players decays in two distinct stages. In theory, the findings could help those who compete for society's continued attention — from politicians and companies to environmental campaigners — to find ways to stay in the public eye, or at least in the public's head.

The study applies maths and a big-data approach to questions that have been studied at length in the social sciences. Using attention as a proxy for memory, the authors analysed online views of the Wikipedia profiles of around 1,700 sports stars, citations of almost 500,000 physics papers and 1.7 million patents, and online play counts of some 33,000 songs and 15,000 film trailers.

Researchers had previously thought that the decline in the popularity of such cultural objects followed a smooth, steep curve. But analysis of the new study data revealed that a better fit was a shape called a biexponential function, which has two phases. It shows that collective memory dropped quickly, but that the subsequent decline in attention slowed considerably, and went down a much gentler slope. Although the shape was the same for each feature studied, the actual length of each phase was different. Music showed the shortest and sharpest initial decline in attention (taking 6 years) and the online biographies of the sports stars the longest (20–30 years).

How come? The researchers propose an explanation. The first, steep decline phase is dominated by the process of communicative memory, which is the direct word-of-mouth transfer of information. And the second, more enduring phase relies more on cultural memory, which is sustained by the physical recording of that same information.

That requires, of course, that the information is recorded. As an accompanying News & Views article highlights, for events that are memorialized with few cultural artefacts, such as Hurricane Sandy striking New York in 2012, policymakers could look at how to extend the period for which communicative memory dominates (A. Coman *Nature Hum. Behav.* http://doi.org/cxst; 2018). For a short time, conversations about the damage it caused probably raised awareness of climate change as a serious threat. But as collective memory of the severity of the hurricane faded, so, too, did concern.

The model does not apply in all cases, of course. Everyone will have their own example of an enduring figure still waiting in Eagleman's purgatorial lobby for their name to become redundant. But it's a neat way to apply the promise of big data to a new field of study, and one that could have real-world applications. It's also another example of how what can seem to be random and individual events when studied at a large enough scale can reveal an underlying pattern. The researchers compare their biexponential function of collective-memory decay to the more poetic description of a two-phase system from Chilean writer Pablo Neruda: "Love is so short, forgetting is so long." Which, at the very least, should keep Neruda hanging around for a bit longer. ■