When he was still in primary school, Krishna Gummadi learned to play musical instruments and studied programming. He soon gave up on music, but programming turned out to be his calling. These days, as director at the Max Planck Institute for Software Systems in Saarbrücken, he is researching, among other things, why artificial intelligence often makes decisions that are just as discriminatory as the ones humans make, and how this can be prevented.

TEXT: TIM SCHRÖDER

When the first algorithms that came close to matching human intelligence appeared on the market in the early 1990s, they sparked great enthusiasm. Banks used them to take on the time-consuming task of deciphering handwriting on cheques, while others recognized objects in pictures for the first time – soccer balls on grass, for example. Computers were no longer simply processing instruments, mathematical robots that, like a chess program, simply played through thousands of variations in a matter of seconds. Now they could actually recognize and interpret things. Since then, algorithms have begun to take a lot of decisions away from humans, and this is often controversial: they filter out the best candidates from applications for a new position. Other algorithms sense the preferences of Internet shoppers in order to place targeted advertisements. Artificial intelligence (AI) is thus encroaching more deeply than ever into our everyday lives, and our society. "Artificial intelligence has spawned socio-technical systems that have a significant impact on how we live together," says Krishna Gummadi. "What interests me is the problems this brings and how we can solve them."

Krishna Gummadi is the director of the Max Planck Institute for Software Systems in Saarbrücken. He has focused for many years on distributed computing networks, cloud computing, and secure data traffic on the Internet, and for some time now, he has been particularly interested in the merging of society and technology. He calls this "social computing". The extent to which the decisions of "socio-technical systems" are unjust and can disadvantage people is a topic of increasingly frequent debate – including among the public and in the media. Krishna Gummadi examines these algorithms closely.

A few years ago, the public learned about the AI software COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) from the USA, which was supposed to reliably calculate the recidivism risk of offenders. To evaluate this risk, the software used not only information about previous convictions and the severity of current offense, but also personal data such as the age of the offender. Though the software designers denied using additional data, the program also accessed criminal records of close relatives, information about alcohol and drug abuse within the family, social ties, friends, and the person's financial situation. This data was further supplemented by character traits such as tendency toward anger and aggression. In many states, judges handed down particularly harsh sentences based on poor COMPAS scores. Experts from the research network ProPublica examined the COMPAS results more closely – and then published a study that made headlines. It showed that the COMPAS algorithms gave defendants of color a higher risk of re-

→

VISIT TO

KRISHNA GUMMADI

PHOTO: ANNA ZIEGLER FOR MPG

Cross-border researcher: Krishna Gummadi carries out research at the interface between computer science and social sciences. For example, he investigates the social impacts of artificial intelligence.

PHOTO: ANNA ZIEGLER FOR MPG

Programmed for fairness: sometimes artificial neural networks make discriminatory decisions. Krishna Gummadi is expanding algorithms to prevent this in the future.

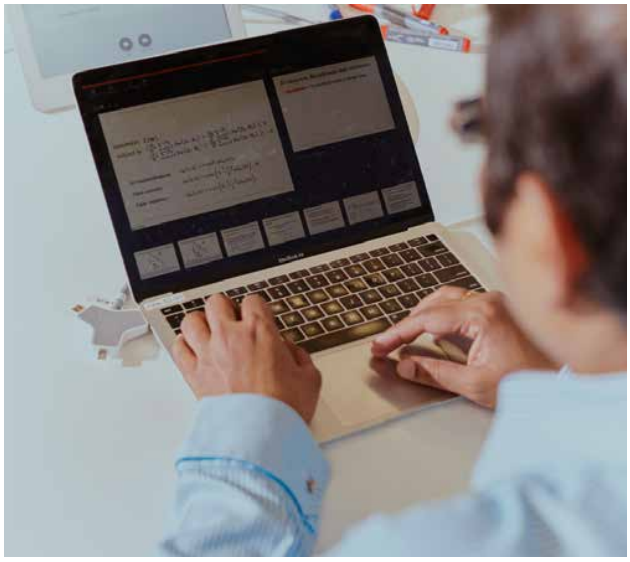

PHOTO: ANNA ZIEGLER FOR MPG

offending than was actually the case. The reverse was true for white defendants: COMPAS gave overly positive predictions more often. AI algorithms like the ones used for COMPAS are based on processes of machine learning.

The algorithms are constantly fed with data from reality and trained to then make decisions independently. Yet these decisions are only ever as good and accurate as the data with which the algorithms are trained. A well-known example is an algorithm that learned to recognize soccer balls in pictures. An analysis of the software showed that the algorithm identified soccer balls by the crIteria "black and white," "hexagonal", and "green" because many of the photos showed grass – a correlation that has nothing to do with the characteristics of soccer balls. The soccer ball algorithm makes for a good anecdote, but the COMPAS algorithm had dire consequences because it discriminated against people of color. Ultimately, COMPAS had been trained with data that had been gathered by human beings – and quite obviously, these were so prejudiced that they disadvantaged such people. "We call this kind of thing a bias, a distortion of the data," says Krishna Gummadi. "Basically, the developers of COMPAS meant well. They wanted to make risk assessment more objective by letting the computer do the work." Humans can be biased, but the computer is not – or so they thought.

There was a similar situation a few years ago with a US software program that was supposed to automatically select suitable employees from a large number of applicants. It was found that it suggested women as suitable candidates significantly less often. "We want to understand how the algorithms work so that these kinds of weaknesses can be ironed out," says Krishna Gummadi. When asked if he is concerned with discrimination because he may have been a victim of it himself, he shakes his head. "No, I'm just interested in a lot of different topics – combining social aspects and computer science is something I really enjoy." Krishna Gummadi often smiles when he discusses his work. "I really enjoy working and actually don't have all that many hobbies."

Krishna Gummadi grew up in the Indian city of Hyderabad. It was important to Gummadi's father that his two sons receive a good education. He had been the first in the family to go to university, completing a bachelor's degree in engineering. At the time, he only had one option: after his bachelor's he had to earn money. Krishna and his brother, meanwhile, were to go further in life. "Our parents made sure we were educated as broadly as possible," he recalls. "Our father enrolled us in various courses when we were still in primary school – guitar lessons, for example, as well as flute lessons and a computer science course, where we gather our first experience of programming." It wasn't long before he gave up the flute and the guitar, but the computer science grabbed him. After finishing school, he – along with around 200,000 other Indian high school graduates in his year – took the university aptitude test, which included questions about various subjects. He came 18th overall. "That was a huge bonus, because the first 20 are free to choose the subject they study." He chose computer science and engineering and moved to Chennai, where he did his bachelor's degree at the Indian Institute of Technology. A master's and doctorate followed at the University of Washington in Seattle, where he worked as a research assistant for several years.

Ending up in Saarbrücken was merely coincidence. "I applied to several universities for tenure in 2005 – including Rice University in Houston. That put me in touch with Peter Druschel, who was working there at the time." Peter Druschel is the founding director of the Max Planck Institute for Software Systems. He was impressed by Gummadi's work and asked him whether he could picture himself moving to Saarbrücken. At first, Gummadi was hesitant about relocating to Germany, but he eventually agreed after Peter Druschel offered him a tenure track position – a permanent role that can lead to a professorship. At the time, no one could have foreseen that Krishna Gummadi would one day become director of the Institute – or even that he would stay in Saarbrücken for so long. "But the city is incredibly international, perhaps because of its proximity to France. And thanks to the other computer science institutes on the campus, there's a high density of professional colleagues here." Together with his wife, he lives on the Saar in a little village, which is next to the border with France. And there's another thing he particularly likes – the region's cuisine. As things stand, Krishna Gummadi is here to stay.

And that's despite the fact that even after so many years in Germany, he still prefers to conduct longer conversations in English. "The Institute is so international that everyone speaks English here, which means you don't practice much on a day-to-day basis." He speaks German, of course, when he does his shopping, but because his wife is also from India, there is no reason to speak it at home. When a new

→

doctoral researcher recently asked if it was possible to survive in Germany without speaking German, the people in his working group laughed. "I'm the best example of that," they said.

For the past five years or so, Krishna Gummadi has been getting more and more involved with "social computing". With the development of AI algorithms that use machine learning, you give the computer a goal and then just let the algorithm do it. "We call this the declarative approach, where I just define the goal – for example, 'Pick the best applicant.' The path through the individual development steps is of no concern," says Krishna Gummadi. The tools used in this process include so-called deep neural networks. Although these algorithms and neural networks deliver results, it is no longer possible to see how the computer made its decision. The process is like a black box, and this becomes problematic, among other things, when women are less likely to be invited to a job interview because the computer filters them out beforehand.

"Just a few years ago, computer scientists doubted that such cases were computer engineering problems," says Krishna Gummadi. "Now it's clear that with socio-technical systems, information science is one of the tools we need to solve problems." Instead of pursuing just one goal in a "utilitarian" fashion as in the past – aiming for an outcome with the lowest possible error rate – in the future, additional goals must be defined, such as preventing unequal treatment and discrimination. The challenge today is to also teach the computer the social context.

One of the definitions of discrimination is: "Discrimination is the act of distinguishing between people on the basis of the groups, classes, or other categories to which they belong or are perceived to belong." Computer scientists must cast references such as "based on groups" or "belonging" into algorithms – an abstraction task that initially has nothing to do with bits, ones, and zeros. "We spent a long time thinking about how you can abstract discrimination in order to develop algorithms that will be free of discrimination in the future. In the process, we came up with the notion of envy-free," says Krishna Gummadi. Among other things, membership in groups becomes problematic when one is favored and the other disadvantaged, which can generate envy, he explains. And that can be expressed mathematically.

"Now it's clear that with socio-technical systems, information science is one of the tools we need to solve problems."

Krishna Gummadi's social computing goes beyond the topic of discrimination. He is interested in how algorithms work in "digital public spaces" – on platforms such as Facebook and TikTok or at online mail-order companies such as Amazon. In a recent study, he and colleagues at the Indian Institute of Technology examined the extent to which Amazon's website algorithms discriminate against vendors. Amazon has long been more than just a mail order company. They also manufacture their own products – and offer them in competition with established manufacturers. Gummadi's team examined how often on the Amazon site the window "Other customers also bought..." displayed Amazon's own products or the goods of other manufacturers. They focused on batteries and backpacks, and the results are sobering: in the case of backpacks, Amazon displayed its own products twice as often. This is admittedly not illegal, but Amazon has another major competitive advantage: the company simultaneously acquires countless customer and market data, which are not available to the other providers. "In the long run, this can weaken the market position of the other providers and is basically a case for the regulators," says

Krishna Gummadi. "Amazon can't be both a player and the referee setting the rules." He also believes it is his duty to point out such abuses in a scientifically sound manner, but says it's up to others to take action.

Back in 2015, he and Emilio Zagheni, director at the Max Planck Institute for Demographic Research in Rostock, had already investigated what data Facebook manages in the so-called API, a programming interface, as part of a project together. Via the API, companies can place advertisements and leverage the browsing behavior of users and potentially interested parties. A whole universe of new data opened up for social science research and particularly, for demographic research. The researchers were able to study migration patterns, including migration of refugees during crises. When conducting the studies, Gummadi's team discovered a privacy vulnerability due to a bug in the API design. A malicious advertiser would be able to retrieve an explosive amount of people's private data, including their addresses, phone numbers, and all the personal information that came from databases to which Facebook is linked. "The amount of private data that was exposed here was staggering," says Krishna Gummadi. "Through the API, we had access to several thousand attributes." Together with his collaborators, Gummadi published a specialist article on the matter, and Facebook has since revised the API interfaces.

"It's funny," he says. "Social computing is a hot topic today – but bringing together social and technical aspects is basically old hat." For example, he says, many of the great computer scientists of the early days had a background in the social sciences – and only later tackled topics such as cognition or decision-making, paving the way for artificial intelligence. In this respect, he says, social computing has now come full circle. Two years ago, Krishna Gummadi organized a symposium for the Max Planck Society on the intersection of society and computer science. Interest was enormous. In just a short time, 270 participants from various Max Planck Institutes signed up.

And the volume of topics is also enormous. Krishna Gummadi and his colleagues are constantly coming across new aspects, but right now he is interested in how the algorithm of the social media platform TikTok works. TikTok delivers an endless stream of short videos. Without people clicking on videos or entering search terms, the platform learns about users' preferences in just a short space of time. How quickly does someone keep scrolling? How long does someone spend watching a video? After only about half an hour, the algorithm delivers videos that fit perfectly – and thus keeps users on the platform. "This leads us to yet other aspects of social computing," says Krishna Gummadi, "to the question of the extent to which the direct, often emotional appeal leads to addictive behavior, to depression or to loneliness." The beauty of working at the Saarbrücken Institute is the freedom to do research without having to deliver results right away, he says. "Here, my curiosity is given free rein – so I can take a detour now and again in subject matter, too."

←

Bound by algorithms: social media platforms like the video platform TikTok are often programmed to cast a spell over users – the psychological consequences of this are still unknown.

PHOTO: ADOBESTOCK / NATTAKORN