nature

Facial-recognition research needs an ethical reckoning

The fields of computer science and AI are struggling with the ethical challenges of biometrics. Institutions must respond.

Over the past 18 months, a number of universities and companies have been removing online data sets containing thousands — or even millions — of photographs of faces used to improve facial-recognition algorithms. In most cases, researchers scraped these images from the Internet. The pictures were public, and their collection didn't seem to alarm institutional review boards (IRBs) and other research-ethics bodies. But none of the people in the photos had been asked for permission, and some were unhappy about the way their faces had been used.

This problem has been brought to prominence by the work of Berlin-based artist and researcher Adam Harvey, who highlighted how public data sets are used by companies to hone surveillance-linked software — and by the journalists who reported on Harvey's work. Many researchers in the fields of computer science and artificial intelligence (AI), and those responsible for the relevant institutional ethical review processes, did not see any harm in using public data without consent. But that is starting to change. It is one of many debates that need to be had around how facial-recognition work — and many other kinds of AI research — can be studied more responsibly.

As *Nature* reports in a series of Features on facial recognition this week (pages 347, 350 and 354), many in the field are rightly worried about how the technology is being used. Some scientists are analysing inaccuracies and biases in facial-recognition technology, warning of discrimination, and joining campaigners calling for stronger regulation, greater transparency, consultation with the communities that are being monitored by cameras — and for use of the technology to be suspended while lawmakers reconsider its benefits and risks.

Responsible studies

Some scientists are urging a rethink of ethics in the field of facial-recognition research, too. They are arguing, for example, that researchers should not do certain types of study. Many are angry about academic papers that sought to study the faces of people from vulnerable groups, such as the Uyghur population in China, whom the government has subjected to surveillance and detained on a mass scale.

Others have condemned papers that sought to classify faces by scientifically and ethically dubious measures such as criminality.

> **"Scientists shouldn't gather and analyse personal data simply because they can."**

Nature conducted a survey to better understand scientists' views on the ethics of facial-recognition technology and research. Many respondents said that they wanted conferences to introduce ethics reviews for biometrics studies. This is starting to happen. Next month's NeurIPS (Neural Information Processing Systems) conference will, for the first time, require that scientists address ethical concerns and potential negative outcomes of their work. And the journal *Nature Machine Intelligence* has begun to ask researchers to write a statement describing the impact of certain types of AI research. These are important steps, but journals, funders and institutions could do more.

For example, researchers want more guidance from institutions on what kinds of study are acceptable. *Nature*'s survey found widespread worry — and disagreement — about the ethics of facial-recognition studies, and concern that IRBs might not be equipped to provide sufficient guidance.

General ethical guidance for AI already exists. And US and European funders have supported efforts to study the challenges of biometrics research and recommended rethinking what counts as 'public' data — as well as urging scientists to consider a study's potential impacts on society. Ultimately, biometrics research involves people, and scientists shouldn't gather and analyse personal data simply because they can. Public consultation is key: scientists should consult those whom the data describe. If this is impossible, researchers should try to reach a panel of representatives who can speak for them.

One problem is that AI guidance tends to consist of principles that aren't easily translated into practice. Last year, the philosopher Brent Mittelstadt at the University of Oxford, UK, noted that at least 84 AI ethics initiatives had produced high-level principles on the ethical development and deployment of AI (B. Mittelstadt *Nature Mach. Intell.* **1**, 501–507; 2019). These converged around medical-ethics concepts, such as respect for human autonomy, the prevention of harm, fairness and explicability (or transparency). But Mittelstadt pointed out that different cultures disagree fundamentally on what principles such as 'fairness' or 'respect for autonomy' mean in practice. Medicine has internationally agreed norms for preventing harm to patients, and robust accountability mechanisms. AI lacks these, Mittelstadt noted. Specific case studies and worked examples would be more helpful to prevent ethics guidance becoming little more than window-dressing.

A second concern is that many researchers depend on companies for their funding and data. Although most firms are concerned by ethical questions about the way biometrics technology is studied and used, their views are likely to be conflicted because their bottom line is to sell products.

Researchers alone can't stop companies and governments using facial-recognition technology and analysis tools unethically. But they can argue loudly against it, and campaign for stronger governance and transparency. They can also reflect more deeply on why they're doing their own work; how they've sourced their data sets; whether the community they're expecting their studies to benefit wants the research to be done; and what the potential negative consequences might be.