

IMPACT STATEMENT

Racial and Gender Equity in Technological Infrastructures

Race and gender are identities we all have. Thus, technologies that define, read, classify, and output information about race and gender will impact every one of us. Embedding identity into machine learning systems have widespread implications for how human beings interact with such technological systems and their broader social world. Conceptions of identity have the potential to reify stereotypes, to restrict access to necessary resources, and, at its worst, to enable intentionally unethical and discriminatory practices. Such outcomes affect people of all identities, but are more likely to disadvantage historically marginalized communities (e.g. people of color and trans people). As evidenced by how historical classification systems have been used (e.g. Apartheid, physiognomy), ensuring equity in algorithmic classification technologies like facial analysis is crucial to building a just society. Equity and social justice for all is core to my drive as a researcher.

I believe that more critical analyses of facial analysis technologies will improve their utility and fairness for all. My broader goal is to *mitigate technological harm* through *improved interdisciplinary collaboration* in machine learning. To ensure facial analysis technologies do not harm individuals and communities, it is necessary to understand identity in computer vision infrastructures holistically, as an interconnected system of inputs and outputs, decisions both human and machine. Identifying and understanding how “gender” and “race” are embedded in layers of computer vision infrastructures will contribute to understanding how gender and race can be *used as affordances*—and how and when humans with differing expertise, whether policy, critical theory, or engineering, can intervene.

By evaluating multiple, interconnected levels of human and machine decision-making in facial analysis software, my research will contribute a framework for addressing identity ethically throughout the system development or research lifecycle. An identity framework will directly benefit the design and evaluation of computer vision systems, yet be easily transferable to broader algorithmic infrastructures, for creating more robustly defined and contextually situated models. The desired outcome will be to ensure that identity is handled with care at every stage of the research and development pipelines, inevitably shaping a more equitable and ethical AI and machine learning future.