

Thesis Project Portfolio

Enhancing Minority Business Representation through Machine Learning: A Case Study in Fairfax County

(Technical Report)

White Box vs. Black Box: The Ethical Tradeoffs of AI Transparency

(STS Research Paper)

An Undergraduate Thesis

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Anjali Mehta

Spring, 2025

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

Enhancing Minority Business Representation through Machine Learning: A Case Study in Fairfax County

White Box vs. Black Box: The Ethical Tradeoffs of AI Transparency

Prospectus

Sociotechnical Synthesis

This synthesis connects two interrelated studies: a technical report on identifying minority-owned businesses using machine learning and a STS analysis of algorithmic transparency. The technical work demonstrates dataset biases—revealing that minority-owned businesses in Fairfax County are significantly undercounted (7% in existing data vs. 41.75% in the model). The socio-technical paper cautions that AI tools, especially those using sensitive demographic inferences (e.g., name-based ethnicity classification), risk perpetuating harm if deployed without transparency. Together, these papers highlight a core tension: AI can address systemic inequities but introduces ethical risks if accountability systems are overlooked.

The technical model's design choices mirror key debates from the socio-technical analysis. By using an interpretable decision tree rather than a black-box model, the Fairfax County classifier prioritizes transparency. Yet challenges like small training datasets and unverified external sources underscore the STS paper's warning that even well-intentioned models can exacerbate biases without inclusive data practices. The STS framework also clarifies stakeholder tensions: while policymakers may value the model's economic insights, certain communities might demand greater control over how their demographic data is used and classified.

Achieving the technical report's goal of equitable policymaking requires directly addressing the ethical concerns raised in the STS paper. This includes adopting participatory design practices—such as involving minority business owners in reviewing and validating classification labels—and promoting context-dependent transparency, like explaining model logic without compromising sensitive information. At the same time, the socio-technical critique becomes more impactful when applied to real-world examples like the Fairfax County model. Taken together, the two studies argue for a balanced approach to AI development—one that combines technical innovation with thoughtful governance. This ensures that tools intended to support marginalized communities do not unintentionally perpetuate the very inequalities they aim to address.