

Thesis Project Portfolio

Pairings: Enhancing Food and Wine Experiences Using Machine Learning

(Technical Report)

Examining Algorithmic Bias in Recommender Systems: Ethical Challenges and Solutions

(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

Eli Herman

Spring 2025

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

Pairings: Enhancing Food and Wine Experiences Using Machine Learning

Examining Algorithmic Bias in Recommender Systems: Ethical Challenges and Solutions

Prospectus

Sociotechnical Synthesis

Algorithms now guide many everyday choices, from what we watch to which wine we put on the dinner table. My capstone project, *Pairings*, creates a wine-recommendation app that suggests bottles based on a person's meal, taste, and budget. At the same time, my STS study examines bias in recommendation systems that can squeeze out variety and hide small producers. Treating the two efforts as one story lets me ask a single question: *How can we give people helpful suggestions without funneling them into a narrow groove?*

Pairings tackles a common dining hurdle: most diners are unsure which wine will complement both their dish and their budget. The app uses two models working together—one that matches flavors and one that learns from other users—to produce a short list of wines in real time. After drinking, users rate the pick, and these ratings feed the model so it keeps learning. I tried to prevent a “blockbuster effect” by populating the database with bottles from small vineyards and lesser-known regions. Testers say the app not only improves their meals but also introduces them to wines they would never have tried.

That choice grew from my STS research. Even though wine seems low-stakes, the same feedback loops that drive any recommender can shrink cultural choice and mirror social bias. I studied cases from Netflix and Spotify and used fairness and value-sensitive design as my guides. Three lessons stood out. First, lists that rank by raw popularity keep pushing the popular and burying the rest. Second, when users cannot see why the system chose an item, they stop trusting it and feel like numbers, not people. Third, fixes work best when they touch every stage, from balanced training data to goals that reward diversity to simple user controls.

I put those lessons back into *Pairings*. I run checks to see whether certain grapes, regions, or price tiers start to disappear from the top suggestions. During training, the loss function rewards both predicted user scores and a basic diversity measure. The interface now shows a clear “Why this wine?” link that explains each match in plain language and lets users ask for more or fewer wines of that style.

Looking at both parts together shows why engineers and social analysts need each other. The

STS lens warned me about hidden risks, and writing actual code showed how broad ethical ideas become data tables, model weights, and screen text. The finished tool still gives quick, accurate help but also protects user choice and gives small producers a fair chance. This experience suggests that caring about social context early in the build process leads to stronger and fairer technology—by weaving ethics into every design decision, not tacking it on at the end.