

# Pairings: Enhancing Food and Wine Experiences Using Machine Learning

CS4991 Capstone Report, 2024

Eli Herman  
Computer Science  
The University of Virginia  
School of Engineering and Applied Science  
Charlottesville, Virginia USA  
hzt7wt@virginia.edu

## ABSTRACT

With the increased demand for personalized dining experiences, the ability to pair food and wine based on individual preferences has become an essential part of fine dining. I addressed the need for a smart wine pairing solution by developing a machine learning-based app, *Pairings* which suggests wines based on meal components, user taste preferences, and budget. The app utilizes a comprehensive wine and food database, along with sommelier knowledge, to generate recommendations tailored to each user's unique dining context. Through the integration of user feedback and sommelier ratings, the app refines suggestions over time, ensuring continuous improvement in the quality of recommendations. Initial results indicate a high degree of user satisfaction, with personalized wine pairings enhancing the overall dining experience. Future developments will focus on expanding the wine database, incorporating advanced deep learning techniques to further enhance the recommendation engine, and improving the app's user interface for wider adoption.

## 1. INTRODUCTION

In the modern dining experience, personalized recommendations for food and wine pairings are becoming increasingly important. The integration of technology into gastronomy has opened new possibilities for offering highly tailored dining experiences based on

individual preferences. With the global wine industry expected to grow significantly in the future, restaurants, sommeliers, and consumers alike are searching for innovative solutions that enhance their dining experiences. I addressed this issue by developing a machine learning-driven app, *Pairings*, which intelligently suggests wine pairings based on a user's meal, preferences and budget. By leveraging a combination of a wine and food database, sommelier expertise, and user feedback, the app refines its recommendations to create an increasingly personalized experience.

Simultaneously, the societal concerns about bias in algorithmic decision-making, especially in the context of recommender systems, raise questions about fairness, transparency and diversity of choice. In my STS research, I explored the potential biases inherent in machine learning algorithms that could limit exposure to lesser-known wines and reinforce popular or more mainstream choices. Investigating both the technical implementation and the sociotechnical implications of such systems is crucial to understanding their long-term impact on the food and beverage industry. I address the technical aspects of wine pairing algorithms and examine the broader question of algorithmic bias within recommender systems, using *Pairings* as a case study.

## 2. RELATED WORKS

The development of *Pairings* builds on existing research in both recommender systems and sommelier knowledge. Early work by Resnick and Varian (1997) set the stage for the development of modern recommender systems by exploring collaborative filtering techniques for personalized recommendations. More recent studies, such as those by Zhang, et al. (2021), have explored how machine learning algorithms can be applied to the domain of food and wine pairing. Their research shows that AI systems can identify complex relationships between food ingredients and wine characteristics, offering new ways to enhance dining experiences.

However, one critical issue with recommender systems is algorithmic bias, which has been extensively studied in various domains. For instance, Mehrabi, et al. (2021) examine how biases in training data can reinforce preexisting inequalities and limit the diversity of options presented to users. Applying these findings to the wine industry, there is potential for algorithms to recommend popular or expensive wines, thereby neglecting smaller producers or less mainstream options. This concern aligns with the work of Ekstrand, et al. (2018), who emphasize the importance of fairness in recommender systems, particularly in fields like entertainment and e-commerce, but with clear implications for food and wine recommendations.

By integrating machine learning into a domain traditionally dominated by human expertise, I draw from both technical research on recommender systems and sociotechnical discussions on algorithmic bias. Through the development of *Pairings*, I demonstrate that machine learning can offer valuable insights in the culinary world while acknowledging the risks associated with biased recommendations.

## 3. PROJECT DESIGN

The *Pairings* app is an application of machine learning in the niche domain of food and wine pairing. This section provides an overview of the app's design, detailing its architecture, data flow, machine learning model, and mechanisms for user feedback integration, as well as addressing the technical challenges encountered during development.

### 3.1 Overview

*Pairings* was developed as a hybrid recommendation system that integrates content-based filtering with collaborative filtering to suggest wines based on meal components, user preferences and budget. Such hybrid approaches are known to improve recommendation accuracy by combining multiple perspectives (Fernández-Tobías et al., 2011). This design utilizes a rich dataset of wine profiles and sommelier knowledge to create tailored recommendations. Through continuous integration of user feedback and ratings, the app learns and refines its suggestions, ensuring ongoing improvement in recommendation quality.

### 3.2 Architecture and Data Flow

The architecture is divided into three core components:

- **User Interface:** A web and mobile-friendly interface allows users to input meal details, select taste preferences, and specify a budget. The intuitive design encourages user interaction, similar to best practices identified in consumer-facing applications (Verbert, et al., 2012).
- **Backend Server:** The backend, developed in Python with TensorFlow, processes user inputs through machine learning models. It integrates the recommendation algorithms, ensuring real-time responses and personalized suggestions.
- **Database:** The database holds extensive data on wine characteristics, including varietals, regions, and sommelier ratings.

User interactions are also logged, supporting the iterative improvement of recommendations. Data structuring follows guidelines for effective knowledge management in recommendation systems (Resnick & Varian, 1997).

### 3.3 Machine Learning Model

The app employs a hybrid recommendation system:

- **Content-Based Filtering:** This aspect matches wines to meals by analyzing flavor profiles and known pairing principles, similar to approaches in other niche recommendation systems (Verbert et al., 2012).
- **Collaborative Filtering:** It uses user behavior data to find patterns in preferences, recommending wines enjoyed by users with similar tastes. Such methods have been shown to enhance user satisfaction in personalized applications (Resnick & Varian, 1997).
- **Model Training:** Training data includes user feedback, meal components, and wine attributes, allowing the system to identify which pairings users appreciate the most. Over time, this improves the accuracy of recommendations, as seen in cross-domain recommendation systems (Fernández-Tobías et al., 2011).

### 3.4 User Feedback Integration

User feedback plays a critical role in refining the model. After receiving a recommendation, users rate their experience, which feeds into the system's learning process. This feedback loop not only adjusts the model's suggestions but also ensures that the algorithm remains aligned with user preferences. Such adaptive learning mechanisms have been found to significantly enhance recommendation quality (Verbert et al., 2012).

### 3.5 Technical Challenges

Developing *Pairings* posed challenges such as balancing diverse data sources and addressing algorithmic bias. As studies like Noble (2018)

and O'Neil (2016) have highlighted, bias in recommendation systems can reinforce pre-existing preferences and exclude lesser-known options. *Pairings* mitigates this by using diverse datasets and transparency-focused algorithms, aiming to provide fair and varied recommendations.

## 4. ANTICIPATED RESULTS

The anticipated outcome of my technical project, the *Pairings* app, is a robust recommendation system that significantly enhances user dining experiences by delivering personalized wine suggestions tailored to individual meal components, taste preferences, and budgets. By combining content-based and collaborative filtering techniques, I aim to provide recommendations that adapt over time based on user feedback and emerging wine trends. User satisfaction is anticipated to improve as the model continues to refine its suggestions, building a strong base of customer loyalty and engagement. Additionally, this app has the potential to disrupt traditional wine pairings, making high-quality recommendations accessible to a broader audience.

On the STS side, my research aims to contribute a critical analysis of algorithmic bias in recommender systems, especially in niche areas like wine selection. This research is expected to highlight how recommender algorithms can unintentionally limit diversity by favoring popular wines and regions, potentially stifling lesser-known options and consumer choice. The findings should offer actionable insights for mitigating bias in recommendation systems, such as incorporating diverse datasets and improving transparency in algorithmic design. Together, these outcomes could pave the way for creating more inclusive and equitable recommendation algorithms across various industries, fostering a broader cultural

appreciation for diverse wine options and offering fairer choices for consumers.

## 5. CONCLUSION

The *Pairings* app represents an innovative step in integrating machine learning with the dining experience. By addressing a clear need for personalized and accessible wine recommendations, it leverages advanced algorithms and diverse datasets to provide tailored suggestions based on meal components, preferences, and budgets.

Beyond enhancing individual dining experiences, the project underscores the potential of recommender systems to bring transparency and diversity to traditionally subjective areas like wine pairing. This research also deepens our understanding of the impact of algorithmic decision-making in niche markets, aligning technical achievements with broader societal needs.

## 6. FUTURE WORK

Moving forward, the *Pairings* project will focus on expanding its dataset to include a wider variety of wines, food profiles, and user feedback to refine its accuracy. Advanced techniques such as deep learning could be explored to uncover more complex pairing patterns. The app's user interface will undergo iterative design improvements to enhance accessibility and engagement.

Additionally, future iterations will aim to include real-time feedback integration and broader language and cultural customizations, enabling the app to cater to a global audience. Beyond wine, the app's framework could be adapted for other domains, such as cocktail or craft beer pairings, extending its utility and market reach.

## REFERENCES

- Ekstrand, M. D., Tian, M., Azpiazu, J., Ekstrand, J. D., Anuyah, O., McNeill, D.,

... & Pera, M. S. (2018). All the cool kids, how do they fit in?: Popularity and demographic biases in recommender evaluation and effectiveness. *Proceedings of the 12th ACM Conference on Recommender Systems*.

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6), 1-35.

Resnick, P., & Varian, H. R. (1997). Recommender systems. *Communications of the ACM*, 40(3), 56-58.

Zhang, Z., Li, S., & Jafari, A. (2021). Machine learning approaches for food-wine pairing: From classical learning to deep learning. *Journal of Food Engineering*, 298, 110469.

Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachsler, H., Bosnic, I., & Duval, E. (2012). Context-aware recommender systems: A review and future directions. *User Modeling and User-Adapted Interaction*, 22(1-2), 157-190.

Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.

Fernández-Tobías, I., Cantador, I., Kaminskas, M., & Ricci, F. (2011). A generic semantic-based framework for cross-domain recommendation. *Proceedings of the Fifth ACM Conference on Recommender Systems (RecSys '11)*, 33-40.